The Internet of Things: An Overview

Understanding the Issues and Challenges
of a More Connected World

OCTOBER 2015

We organize this paper into three main sections:

- **What is the Internet of Things?**, which provides an overview of its origins, definitions, and technical connectivity models;

- **What issues are raised by the Internet of Things?**, which provides an introduction and discussion of concerns that have been raised about IoT, and;

- **For Further Information,** which provides additional information and pointers to efforts around the world addressing IoT issues.

What is the Internet of Things?

Origins, Drivers, and Applications

The term “Internet of Things” (IoT) was first used in 1999 by British technology pioneer Kevin Ashton to describe a system in which objects in the physical world could be connected to the Internet by sensors.[12] Ashton coined the term to illustrate the power of connecting Radio-Frequency Identification (RFID) tags[13] used in corporate supply chains to the Internet in order to count and track goods without the need for human intervention. Today, the Internet of Things has become a popular term for describing scenarios in which Internet connectivity and computing capability extend to a variety of objects, devices, sensors, and everyday items.

While the term “Internet of Things” is relatively new, the concept of combining computers and networks to monitor and control devices has been around for decades. By the late 1970s, for example, systems for remotely monitoring meters on the electrical grid via telephone lines were already in commercial use.[14] In the 1990s, advances in wireless technology allowed “machine–to–machine” (M2M) enterprise and industrial solutions for equipment monitoring and operation to become widespread. Many of these early M2M solutions, however, were based on closed purpose–built networks and proprietary or industry–specific standards,[15] rather than on Internet Protocol (IP)–based networks and Internet standards.

Using IP to connect devices other than computers to the Internet is not a new idea. The first Internet “device”—an IP–enabled toaster that could be turned on and off over the Internet—was featured at an Internet conference in 1990.[16] Over the next several years, other “things” were IP–enabled, including a soda machine[17] at Carnegie Mellon University in the US and a coffee pot[18] in the Trojan Room at the University of Cambridge in the UK (which remained Internet–connected until 2001). From these whimsical beginnings, a robust field of research and development into “smart object networking”[19] helped create the foundation for today’s Internet of Things.

[12] Ashton was working on RFID (radio-frequency identification) devices, and the close association of RFID and other sensor networks with the development of the IoT concept is reflected in the name of the RFID device company that Ashton joined later in his career: “ThingMagic.”

[13] “Radio-Frequency Identification.” *Wikipedia, the Free Encyclopedia*, September 6, 2015. https://en.wikipedia.org/wiki/Radio-frequency_identification

[14] “Machine to Machine.” *Wikipedia, the Free Encyclopedia*, August 20, 2015. https://en.wikipedia.org/wiki/Machine_to_machine

[15] Polsonetti, Chantal. “Know the Difference Between IoT and M2M.” *Automation World*, July 15, 2014. http://www.automationworld.com/cloud-computing/know-difference-between-iot-and-m2m

[16] "The Internet Toaster." Living Internet, 7 Jan. 2000. Web. 06 Sept. 2015. http://www.livinginternet.com/i/ia_myths_toast.htm

[17] "The "Only" Coke Machine on the Internet." Carnegie Mellon University Computer Science Department, n.d. Web. 06 Sept. 2015. https://www.cs.cmu.edu/~coke/history_long.txt

[18] Stafford-Fraser, Quentin. "The Trojan Room Coffee Pot." N.p., May 1995. Web. 06 Sept. 2015. http://www.cl.cam.ac.uk/coffee/qsf/coffee.html

[19] RFC 7452, “Architectural Considerations in Smart Object Networking” (March 2015), https://tools.ietf.org/html/rfc7452

If the idea of connecting objects to each other and to the Internet is not new, it is reasonable to ask, "Why is the Internet of Things a newly popular topic today?"

From a broad perspective, the confluence of several technology and market trends[20] is making it possible to interconnect more and smaller devices cheaply and easily:

- *Ubiquitous Connectivity*—Low–cost, high–speed, pervasive network connectivity, especially through licensed and unlicensed wireless services and technology, makes almost everything "connectable''.
- *Widespread adoption of IP–based networking*— IP has become the dominant global standard for networking, providing a well–defined and widely implemented platform of software and tools that can be incorporated into a broad range of devices easily and inexpensively.
- *Computing Economics*— Driven by industry investment in research, development, and manufacturing, Moore's law[21] continues to deliver greater computing power at lower price points and lower power consumption.[22]
- *Miniaturization*— Manufacturing advances allow cutting-edge computing and communications technology to be incorporated into very small objects.[23] Coupled with greater computing economics, this has fueled the advancement of small and inexpensive sensor devices, which drive many IoT applications.
- *Advances in Data Analytics*— New algorithms and rapid increases in computing power, data storage, and cloud services enable the aggregation, correlation, and analysis of vast quantities of data; these large and dynamic datasets provide new opportunities for extracting information and knowledge.
- *Rise of Cloud Computing*– Cloud computing, which leverages remote, networked computing resources to process, manage, and store data, allows small and distributed devices to interact with powerful back-end analytic and control capabilities.

From this perspective, the IoT represents the convergence of a variety of computing and connectivity trends that have been evolving for many decades. At present, a wide range of industry sectors – including automotive, healthcare, manufacturing, home and consumer electronics, and well beyond -- are considering the potential for incorporating IoT technology into their products, services, and operations.

[20] Other views on the converging market trends driving IoT's growth include Susan Conant's article "The IoT will be as fundamental as the Internet itself", available at http://radar.oreilly.com/2015/06/the-iot-will-be-as-fundamental-as-the-internet-itself.html and Intel Corporation's statement to U.S. House of Representatives hearing on IoT, available at http://docs.house.gov/meetings/IF/IF17/20150324/103226/HHRG-114-IF17-Wstate-SchoolerR-20150324.pdf.

[21] Moore's Law is named after a trend observed by semiconductor pioneer Gordon Moore that the number of transistors per square inch on integrated circuits doubles roughly every two years, allowing more processing power to be placed into smaller chips over time.

[22] For a discussion about Internet device energy use and low power computing, see the lecture by Jon Koomey at the "How green is the Internet?" summit available at https://www.youtube.com/embed/O8-LDLyKaBM

[23] In addition to other technical advancements, miniaturization of electronic devices is also fueled by Moore's law.

In their report “Unlocking the Potential of the Internet of Things’’, the McKinsey Global Institute[24] describes the broad range of potential applications in terms of “settings” where IoT is expected to create value for industry and users.

“Settings” for IoT Applications (Source: McKinsey Global Institute[25])

Setting	Description	Examples
Human	Devices attached or inside the human body	Devices (wearables and ingestibles) to monitor and maintain human health and wellness; disease management, increased fitness, higher productivity
Home	Buildings where people live	Home controllers and security systems
Retail Environments	Spaces where consumers engage in commerce	Stores, banks, restaurants, arenas – anywhere consumers consider and buy; self-checkout, in-store offers, inventory optimization
Offices	Spaces where knowledge workers work	Energy management and security in office buildings; improved productivity, including for mobile employees
Factories	Standardized production environments	Places with repetitive work routines, including hospitals and farms; operating efficiencies, optimizing equipment use and inventory
Worksites	Custom production environments	Mining, oil and gas, construction; operating efficiencies, predictive maintenance, health and safety
Vehicles	Systems inside moving vehicles	Vehicles including cars, trucks, ships, aircraft, and trains; condition-based maintenance, usage-based design, pre-sales analytics
Cities	Urban environments	Public spaces and infrastructure in urban settings; adaptive traffic control, smart meters, environmental monitoring, resource management
Outside	Between urban environments (and outside other settings)	Outside uses include railroad tracks, autonomous vehicles (outside urban locations), and flight navigation; real-time routing, connected navigation, shipment tracking

[24] Manyika, James, Michael Chui, Peter Bisson, Jonathan Woetzel, Richard Dobbs, Jacques Bughin, and Dan Aharon. “The Internet of Things: Mapping the Value Beyond the Hype.” McKinsey Global Institute, June 2015. p.3. http://www.mckinsey.com/insights/business_technology/the_internet_of_things_the_value_of_digitizing_the_physical_world

[25] *Ibid.*

Internet of Things Communications Models

From an operational perspective, it is useful to think about how IoT devices connect and communicate in terms of their technical communication models. In March 2015, the Internet Architecture Board (IAB) released a guiding architectural document for networking of smart objects (RFC 7452),[39] which outlines a framework of four common communication models used by IoT devices. The discussion below presents this framework and explains key characteristics of each model in the framework.

Device-to-Device Communications

The device-to-device communication model represents two or more devices that directly connect and communicate between one another, rather than through an intermediary application server. These devices communicate over many types of networks, including IP networks or the Internet. Often, however these devices use protocols like Bluetooth,[40] Z-Wave,[41] or ZigBee[42] to establish direct device-to-device communications, as shown in Figure 1.

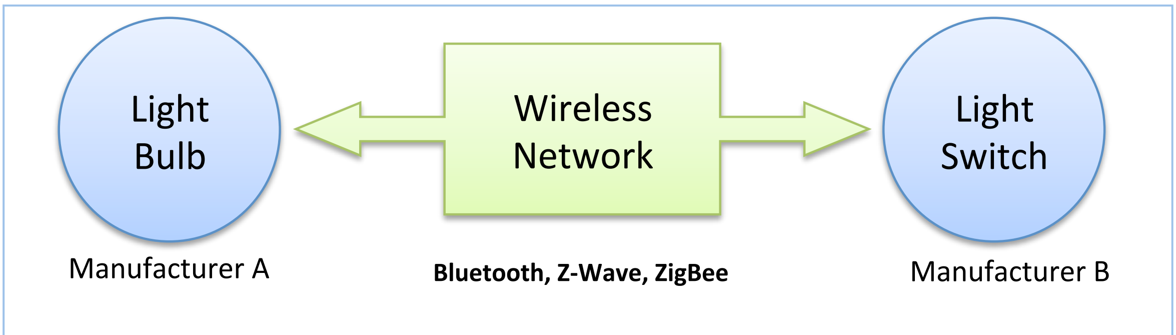

Source: Tschofenig, H., et. al., Architectural Considerations in Smart Object Networking. Tech. no. RFC 7452. Internet Architecture Board, Mar. 2015. Web. <https://www.rfc-editor.org/rfc/rfc7452.txt>.

Figure 1. Example of device-to-device communication model.

These device-to-device networks allow devices that adhere to a particular communication protocol to communicate and exchange messages to achieve their function. This communication model is commonly used in applications like home automation systems, which typically use small data packets of information to communicate between devices with relatively low data rate requirements. Residential IoT devices like light bulbs, light switches, thermostats, and door locks normally send small amounts of information to each other (e.g. a door lock status message or turn on light command) in a home automation scenario.

This device-to-device communication approach illustrates many of the interoperability challenges discussed later in this paper. As an *IETF Journal* article describes, "these devices often have a direct relationship, they usually have built-in security and trust [mechanisms], but they also use device-specific data models that

[39] Tschofenig, H., et. al., *Architectural Considerations in Smart Object Networking*. Tech. no. RFC 7452. Internet Architecture Board, Mar. 2015. Web. https://www.rfc-editor.org/rfc/rfc7452.txt

[40] See http://www.bluetooth.com and http://www.bluetooth.org

[41] See http://www.z-wave.com

[42] See http://www.zigbee.org

require redundant development efforts [by device manufacturers]".[43] This means that the device manufacturers need to invest in development efforts to implement device-specific data formats rather than open approaches that enable use of standard data formats.

From the user's point of view, this often means that underlying device-to-device communication protocols are not compatible, forcing the user to select a family of devices that employ a common protocol. For example, the family of devices using the Z-Wave protocol is not natively compatible with the ZigBee family of devices. While these incompatibilities limit user choice to devices within a particular protocol family, the user benefits from knowing that products within a particular family tend to communicate well.

Device-to-Cloud Communications

In a device-to-cloud communication model, the IoT device connects directly to an Internet cloud service like an application service provider to exchange data and control message traffic. This approach frequently takes advantage of existing communications mechanisms like traditional wired Ethernet or Wi-Fi connections to establish a connection between the device and the IP network, which ultimately connects to the cloud service. This is shown in Figure 2.

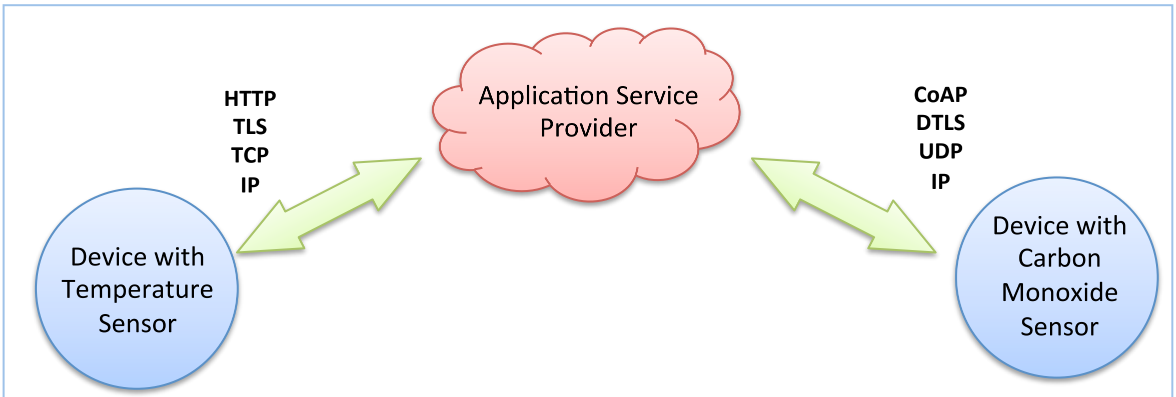
Application Service Provider
HTTP
TLS
TCP
IP
CoAP
DTLS
UDP
IP
Device with Temperature Sensor
Device with Carbon Monoxide Sensor

Source: Tschofenig, H., et. al., Architectural Considerations in Smart Object Networking. Tech. no. RFC 7452. Internet Architecture Board, Mar. 2015. Web. <https://www.rfc-editor.org/rfc/rfc7452.txt>.

Figure 2. Device-to-cloud communication model diagram.

This communication model is employed by some popular consumer IoT devices like the Nest Labs *Learning Thermostat*[44] and the Samsung *SmartTV*.[45] In the case of the Nest *Learning Thermostat*, the device transmits data to a cloud database where the data can be used to analyze home energy consumption. Further, this cloud connection enables the user to obtain remote access to their thermostat via a smartphone or Web interface, and it also supports software updates to the thermostat. Similarly with the Samsung *SmartTV* technology, the television uses an Internet connection to transmit user viewing information to Samsung for analysis and to enable the interactive voice recognition features of the TV. In these cases, the

[43] Duffy Marsan, Carolyn. "IAB Releases Guidelines for Internet-of-Things Developers." *IETF Journal* 11.1 (2015): 6-8. Internet Engineering Task Force, July 2015. Web. https://www.internetsociety.org/sites/default/files/Journal_11.1.pdf

[44] "Meet the Nest Thermostat | Nest." Nest Labs. Web. 31 Aug. 2015. https://nest.com/thermostat/meet-nest-thermostat/

[45] "Samsung Privacy Policy--SmartTV Supplement." Samsung Corp. Web. 29 Sept. 2015. http://www.samsung.com/sg/info/privacy/smarttv.html

device-to-cloud model adds value to the end user by extending the capabilities of the device beyond its native features.

However, interoperability challenges can arise when attempting to integrate devices made by different manufacturers. Frequently, the device and cloud service are from the same vendor.[46] If proprietary data protocols are used between the device and the cloud service, the device owner or user may be tied to a specific cloud service, limiting or preventing the use of alternative service providers. This is commonly referred to as "vendor lock-in'', a term that encompasses other facets of the relationship with the provider such as ownership of and access to the data. At the same time, users can generally have confidence that devices designed for the specific platform can be integrated.

Device-to-Gateway Model

In the device-to-gateway model, or more typically, the device-to-application-layer gateway (ALG) model, the IoT device connects through an ALG service as a conduit to reach a cloud service. In simpler terms, this means that there is application software operating on a local gateway device, which acts as an intermediary between the device and the cloud service and provides security and other functionality such as data or protocol translation. The model is shown in Figure 3.

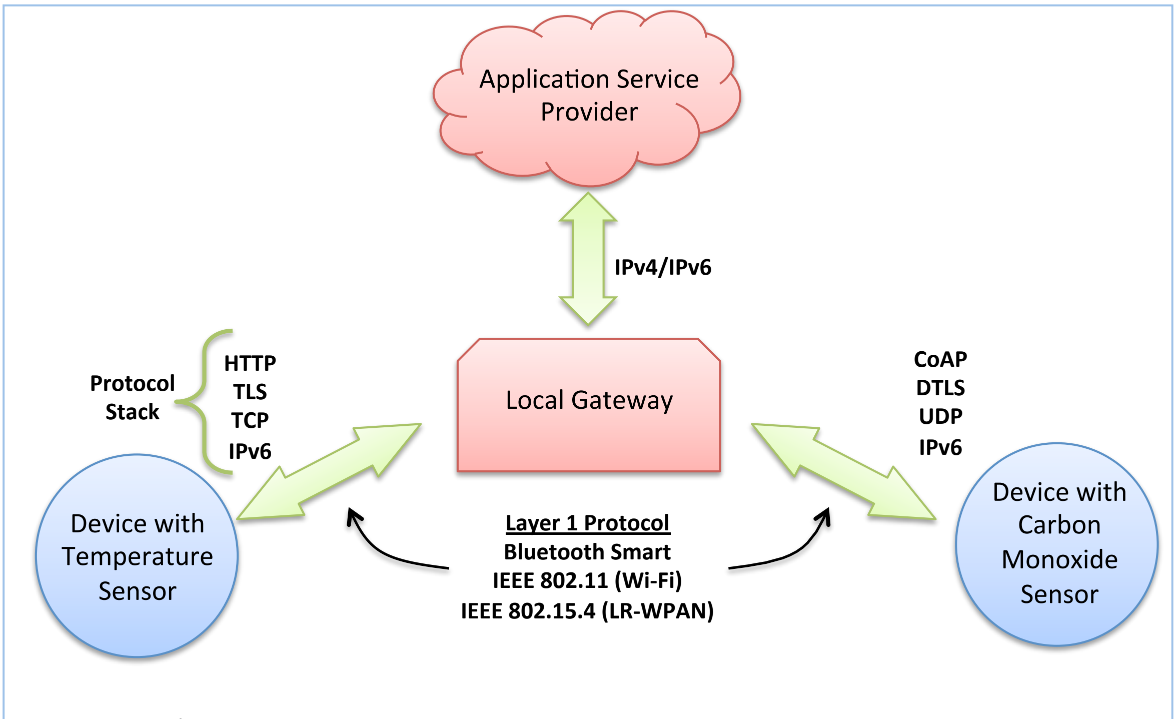

Source: Tschofenig, H., et. al., Architectural Considerations in Smart Object Networking. Tech. no. RFC 7452. Internet Architecture Board, Mar. 2015. Web. <https://www.rfc-editor.org/rfc/rfc7452.txt>.

Figure 3. Device-to-gateway communication model diagram.

Several forms of this model are found in consumer devices. In many cases, the local gateway device is a smartphone running an app to communicate with a device and relay data to a cloud service. This is often the

[46] Duffy Marsan, Carolyn. "IAB Releases Guidelines for Internet-of-Things Developers." *IETF Journal* 11.1 (2015): 6-8. Internet Engineering Task Force, July 2015. Web. https://www.internetsociety.org/sites/default/files/Journal_11.1.pdf

model employed with popular consumer items like personal fitness trackers. These devices do not have the native ability to connect directly to a cloud service, so they frequently rely on smartphone app software to serve as an intermediary gateway to connect the fitness device to the cloud.

The other form of this device-to-gateway model is the emergence of "hub" devices in home automation applications. These are devices that serve as a local gateway between individual IoT devices and a cloud service, but they can also bridge the interoperability gap between devices themselves. For example, the *SmartThings* hub is a stand-alone gateway device that has Z-Wave and Zigbee transceivers installed to communicate with both families of devices.[47] It then connects to the *SmartThings* cloud service, allowing the user to gain access to the devices using a smartphone app and an Internet connection.

From a broader technical perspective, the *IETF Journal* article explains the benefit of the device-to-gateway approach:

> *This [communication model] is used in situations where the smart objects require interoperability with non-IP [Internet protocol] devices. Sometimes this approach is taken for integrating IPv6-only devices, which means a gateway is necessary for legacy IPv4-only devices and services.*[48]

In other words, this communications model is frequently used to integrate new smart devices into a legacy system with devices that are not natively interoperable with them. A downside of this approach is that the necessary development of the application-layer gateway software and system adds complexity and cost to the overall system.

The IAB's RFC7452 document suggests the outlook for this model:

> *It is expected that in the future, more generic gateways will be deployed to lower cost and infrastructure complexity for end consumers, enterprises, and industrial environments. Such generic gateways are more likely to exist if IoT device designs make use of generic Internet protocols and not require application-layer gateways that translate one application-layer protocol to another one. The use of application-layer gateways will, in general, lead to a more fragile deployment, as has been observed in the past…*[49]

The evolution of systems using the device-to-gateway communication model and its larger role in addressing interoperability challenges among IoT devices is still unfolding.

Back-End Data-Sharing Model

The back-end data-sharing model refers to a communication architecture that enables users to export and analyze smart object data from a cloud service in combination with data from other sources. This

[47] "How It Works." *SmartThings*, 2015. http://www.smartthings.com/how-it-works

[48] Duffy Marsan, Carolyn. "IAB Releases Guidelines for Internet-of-Things Developers." *IETF Journal* 11.1 (2015): 6-8. Internet Engineering Task Force, July 2015. Web. https://www.internetsociety.org/sites/default/files/Journal_11.1.pdf

[49] *Tschofenig, H., et. al.,* p. 6.

architecture supports “the [user’s] desire for granting access to the uploaded sensor data to third parties”.[50] This approach is an extension of the single device-to-cloud communication model, which can lead to data silos where “IoT devices upload data only to a single application service provider’’.[51] A back-end sharing architecture allows the data collected from single IoT device data streams to be aggregated and analyzed.

For example, a corporate user in charge of an office complex would be interested in consolidating and analyzing the energy consumption and utilities data produced by all the IoT sensors and Internet-enabled utility systems on the premises. Often in the single device-to-cloud model, the data each IoT sensor or system produces sits in a stand-alone data silo. An effective back-end data sharing architecture would allow the company to easily access and analyze the data in the cloud produced by the whole spectrum of devices in the building. Also, this kind of architecture facilitates data portability needs. Effective back-end data-sharing architectures allow users to move their data when they switch between IoT services, breaking down traditional data silo barriers.

The back-end data-sharing model suggests a federated cloud services approach[52] or cloud applications programmer interfaces (APIs) are needed to achieve interoperability of smart device data hosted in the cloud.[53] A graphical representation of this design is shown in Figure 4.

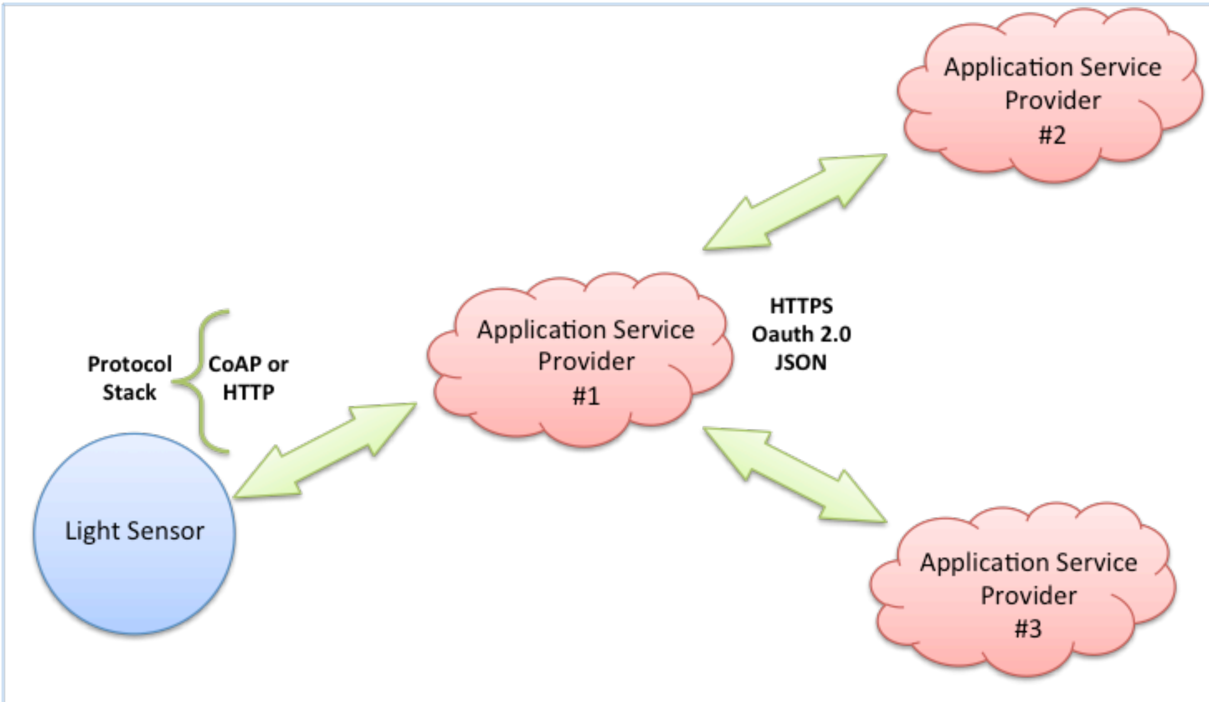
Application Service Provider #2
HTTPS
Oauth 2.0
JSON
Application Service Provider #1
Protocol Stack
CoAP or HTTP
Light Sensor
Application Service Provider #3

Source: Tschofenig, H., et. al., Architectural Considerations in Smart Object Networking. Tech. no. RFC 7452. Internet Architecture Board, Mar. 2015. Web. <https://www.rfc-editor.org/rfc/rfc7452.txt>.

Figure 4. Back-end data sharing model diagram.

[50] *Tschofenig, H., et. al.,* p. 9.

[51] *Ibid.*

[52] A federated cloud services approach is one that combines the resources of separate cloud service providers to meet a larger business need.

[53] An example of a generic (non-IoT) off-the-shelf, federated cloud-sharing tool is *ownCloud,* produced by ownCloud.org. https://owncloud.org/blog/faster-easier-file-sync-and-share-with-federated-self-hosted-owncloud-8-0/

What issues are raised by the Internet of Things?

It would be impossible to cover the broad scope of issues surrounding the Internet of Things in a single paper. Below, however, we provide an overview of five topics frequently discussed in relation to IoT. These include: security; privacy; interoperability and standards; ~~legal, regulatory and rights; and emerging economics and development.~~

We begin to examine these issues through the lens of "the Abilities" – the statement of fundamental principles that guide ISOC's work in terms of the capabilities we believe all Internet users should enjoy that must be protected. These include the ability to *connect, speak, innovate, share, choose*, and *trust*.[55] With these principles as a guide, we present important aspects of each issue and propose several questions for discussion.

Security Issues

The IoT Security Challenge

As we note in the principles that guide our work, ensuring the security, reliability, resilience, and stability of Internet applications and services is critical to promoting *trust* and use of the Internet.[56] As users of the Internet, we need to have a high degree of trust that the Internet, its applications, and the devices linked to it are secure enough to do the kinds of activities we want to do online in relation to the risk tolerance associated with those activities. The Internet of Things is no different in this respect, and security in IoT is fundamentally linked to the ability of users to trust their environment. If people don't believe their connected devices and their information are reasonably secure from misuse or harm, the resulting erosion of trust causes a reluctance to use the Internet. This has global consequences to electronic commerce, technical innovation, free speech, and practically every other aspect of online activities. Indeed, ensuring security in IoT products and services should be considered a top priority for the sector.

As we increasingly connect devices to the Internet, new opportunities to exploit potential security vulnerabilities grow. Poorly secured IoT devices could serve as entry points for cyberattack by allowing malicious individuals to re-program a device or cause it to malfunction. Poorly designed devices can expose user data to theft by leaving data streams inadequately protected. Failing or malfunctioning devices also can create security vulnerabilities. These problems are just as large or larger for the small, cheap, and ubiquitous smart devices in the Internet of Things as they are for the computers that have traditionally been

[55] "Values and Principles." *Principles*. Internet Society, 2015. http://www.internetsociety.org/who-we-are/mission/values-and-principles

[56] *Ibid.*

the endpoints of Internet connectivity. Competitive cost and technical constraints on IoT devices challenge manufacturers to adequately design security features into these devices, potentially creating security and long-term maintainability vulnerabilities greater than their traditional computer counterparts.

Along with potential security design deficiencies, the sheer increase in the number and nature of IoT devices could increase the opportunities of attack. When coupled with the highly interconnected nature of IoT devices, every poorly secured device that is connected online potentially affects the security and resilience of the Internet *globally*, not just locally. For example, an unprotected refrigerator or television in the US that is infected with malware might send thousands of harmful spam emails to recipients worldwide using the owner's home Wi-Fi Internet connection.[57]

To complicate matters, our ability to function in our daily activities without using devices or systems that are Internet-enabled is likely to decrease in a hyperconnected world. In fact, it is increasingly difficult to purchase some devices that are *not* Internet-connected because certain vendors only make connected products. Day by day, we become more connected and dependent on IoT devices for essential services, and we need the devices to be secure, while recognizing that no device can be absolutely secure. This increasing level of dependence on IoT devices and the Internet services they interact with also increases the pathways for wrongdoers to gain access to devices. Perhaps we could unplug our Internet-connected TVs if they get compromised in a cyber attack, but we can't so easily turn off a smart utility power meter or a traffic control system or a person's implanted pacemaker if they fall victim to malicious behavior.

This is why security of IoT devices and services is a major discussion point and should be considered a critical issue. We increasingly depend on these devices for essential services, and their behavior may have global reach and impact.

A Spectrum of Security Considerations

When thinking about Internet of Things devices, it is important to understand that security of these devices is not absolute. IoT device security is not a binary proposition of secure or insecure. Instead, it is useful to conceptualize IoT security as a spectrum of device vulnerability. The spectrum ranges from totally unprotected devices with no security features to highly secure systems with multiple layers of security features. In an endless cat-and-mouse game, new security threats evolve, and device manufacturers and network operators continuously respond to address the new threats.

The overall security and resilience of the Internet of Things is a function of how security risks are assessed and managed. Security of a device is a function of the risk that a device will be compromised, the damage such compromise will cause, and the time and resources required to achieve a certain level of protection. If a user cannot tolerate a high degree of security risk as in the case of the operator of a traffic control system or person with an implanted, Internet-enabled medical device, then she may feel justified in spending a considerable amount of resources to protect the system or device from attack. Likewise, if she is not concerned that her refrigerator might be hacked and used to send spam messages, then she may not feel

[57] Starr, Michelle. "Fridge Caught Sending Spam Emails in Botnet Attack - CNET." CNET, 19 Jan. 2014. http://www.cnet.com/news/fridge-caught-sending-spam-emails-in-botnet-attack/

compelled to pay for a model that has a more sophisticated security design if it makes the device more costly or complicated.

Several factors influence this risk assessment and mitigation calculation. Factors include having a clear understanding of the present security risks and the potential future risks; the estimated economic and other costs of harm if the risks are realized; and the estimated cost to mitigate the risks.[58] While these kinds of security trade-offs are often made from an individual user or organizational perspective, it is also important to consider the interrelatedness of IoT devices as part of a larger IoT ecosystem. The networked connectivity of IoT devices means that security decisions made locally about an IoT device can have global impacts on other devices.

As a matter of principle, developers of smart objects for the Internet of Things have an obligation in ensuring that those devices do not expose either their own users or others to potential harm. As a matter of business and economics, vendors have an interest in reducing their cost, complexity, and time to market. For example, IoT devices that are high–volume, low–margin components that already represent a cost added to that of the product in which they are embedded are becoming quite common; adding more memory and a faster processor to implement security measures could easily make that product commercially uncompetitive.

In economic terms, lack of security for IoT devices results in a negative externality, where a cost is imposed by one party (or parties) on other parties. A classic example is pollution of the environment, where the environmental damage and cleanup costs (negative externalities) of a polluter's actions are borne by other parties. The issue is that the cost of the externality imposed on others is not normally factored into the decision-making process, unless, as is the case with pollution, a tax is imposed on the polluter to convince him to lower the amount of pollution. In the case of information security, as discussed by Bruce Schneier,[59] an externality arises when the vendor creating the product does not bear the costs caused by any insecurity; in this case, liability law can influence vendors to account for the externality and develop more security products.

These security considerations are not new in the context of information technology, but the scale of unique challenges that can arise in IoT implementations, as described below, make them significant.

Unique Security Challenges of IoT Devices

IoT devices tend to differ from traditional computers and computing devices in important ways that challenge security:

- Many Internet of Things devices, such as sensors and consumer items, are designed to be deployed at a massive scale that is orders of magnitude beyond that of traditional Internet-connected devices.

[58] A number of organizations have developed guides for conducting risk assessment. For example, the U.S. National Institute of Standards and Technology (NIST) issued a set of guidelines in 2012, http://www.nist.gov/customcf/get_pdf.cfm?pub_id=912091 and the International Standards Organization (ISO) and the International Electrotechnical Commission (IEC) has published the ISO/IEC 31010:2009 "Risk management – Risk assessment techniques" document. http://www.iso.org/iso/catalogue_detail?csnumber=51073

[59] See Bruce Schneider's online article at: https://www.schneier.com/essays/archives/2007/01/information_security_1.html

As a result, the potential quantity of interconnected links between these devices is unprecedented. Further, many of these devices will be able to establish links and communicate with other devices on their own in an unpredictable and dynamic fashion. Therefore, existing tools, methods, and strategies associated with IoT security may need new consideration.

- Many IoT deployments will consist of collections of identical or near identical devices. This homogeneity magnifies the potential impact of any single security vulnerability by the sheer number of devices that all have the same characteristics. For example, a communication protocol vulnerability of one company's brand of Internet-enabled light bulbs might extend to every make and model of device that uses that same protocol or which shares key design or manufacturing characteristics.

- Many Internet of Things devices will be deployed with an anticipated service life many years longer than is typically associated with high-tech equipment. Further, these devices might be deployed in circumstances that make it difficult or impossible to reconfigure or upgrade them; or these devices might outlive the company that created them, leaving orphaned devices with no means of long-term support. These scenarios illustrate that security mechanisms that are adequate at deployment might not be adequate for the full lifespan of the device as security threats evolve. As such, this may create vulnerabilities that could persist for a long time. This is in contrast to the paradigm of traditional computer systems that are normally upgraded with operating system software updates throughout the life of the computer to address security threats. The long-term support and management of IoT devices is a significant security challenge.

- Many IoT devices are intentionally designed without any ability to be upgraded, or the upgrade process is cumbersome or impractical. For example, consider the 2015 Fiat Chrysler recall of 1.4 million vehicles to fix a vulnerability that allowed an attacker to wirelessly hack into the vehicle. These cars must be taken to a Fiat Chrysler dealer for a manual upgrade, or the owner must perform the upgrade themselves with a USB key. The reality is that a high percentage of these autos probably will *not* be upgraded because the upgrade process presents an inconvenience for owners, leaving them perpetually vulnerable to cybersecurity threats, especially when the automobile appears to be performing well otherwise.

- Many IoT devices operate in a manner where the user has little or no real visibility into the internal workings of the device or the precise data streams they produce. This creates a security vulnerability when a user believes an IoT device is performing certain functions, when in reality it might be performing unwanted functions or collecting more data than the user intends. The device's functions also could change without notice when the manufacturer provides an update, leaving the user vulnerable to whatever changes the manufacturer makes.

- Some IoT devices are likely to be deployed in places where physical security is difficult or impossible to achieve. Attackers may have direct physical access to IoT devices. Anti-tamper features and other design innovations will need to be considered to ensure security.

- Some IoT devices, like many environmental sensors, are designed to be unobtrusively embedded in the environment, where a user does not actively notice the device nor monitor its operating status. Additionally, devices may have no clear way to alert the user when a security problem arises, making it difficult for a user to know that a security breach of an IoT device has occurred. A security breach might persist for a long time before being noticed and corrected if correction or mitigation is even possible or practical. Similarly, the user might not be aware that a sensor exists in her surroundings, potentially allowing a security breach to persist for long periods without detection.

- Early models of Internet of Things assume IoT will be the product of large private and/or public technology enterprises, but in the future "Build Your own Internet of Things" (BYIoT) might become more commonplace as exemplified by the growing *Arduino* and *Raspberry Pi*[60] developer communities. These may or may not apply industry best practice security standards.

IoT Security Questions

A number of questions have been raised regarding security challenges posed by Internet of Things devices. Many of these questions existed prior to the growth of IoT, but they increase in importance due to the scale of deployment of IoT devices. Some prominent questions include:

a) **Good Design Practices.** What are the sets of best practices for engineers and developers to use to design IoT devices to make them more secure? How do lessons learned from Internet of Things security problems get captured and conveyed to development communities to improve future generations of devices? What training and educational resources are available to teach engineers and developers more secure IoT design?

b) **Cost vs. Security Trade-Offs.** How do stakeholders make informed cost-benefit analysis decisions with respect to Internet of Things devices? How do we accurately quantify and assess the security risks? What will motivate device designers and manufacturers to accept additional product design cost to make devices more secure, and, in particular, to take responsibility for the impact of any negative externalities resulting from their security decisions? How will incompatibilities between functionality and usability be reconciled with security? How do we ensure IoT security solutions support opportunities for IoT innovation, social and economic growth?

c) **Standards and Metrics**. What is the role of technical and operational standards for the development and deployment of secure, well-behaving IoT devices? How do we effectively identify and measure characteristics of IoT device security? How do we measure the effectiveness of Internet of Things security initiatives and countermeasures? How do we ensure security best practices are implemented?

d) **Data Confidentiality, Authentication and Access Control.** What is the optimal role of data encryption with respect to IoT devices? Is the use of strong encryption, authentication and access control technologies in IoT devices an adequate solution to prevent eavesdropping and hijacking

[60] See the Arduino open source community http://www.arduino.cc and the Raspberry Pi Foundation http://www.raspberrypi.org/

attacks of the data streams these devices produce? Which encryption and authentication technologies could be adapted for the Internet of Things, and how could they be implemented within an IoT device's constraints on cost, size, and processing speed? What are the foreseeable management issues that must be addressed as a result of IoT-scale cryptography? Are concerns about managing the crypto-key lifecycle and the expected period during which any given algorithm is expected to remain secure being addressed? Are the end-to-end processes adequately secure and simple enough for typical consumers to use?

e) **Field-Upgradeability.** With an extended service life expected for many IoT devices, should devices be designed for maintainability and upgradeability in the field to adapt to evolving security threats? New software and parameter settings could be installed in a fielded IoT device by a centralized security management system if each device had an integrated device management agent. But management systems add cost and complexity; could other approaches to upgrading device software be more compatible with widespread use of IoT devices? Are there any classes of IoT devices that are low-risk and therefore don't warrant these kinds of features? In general, are the user interfaces IoT devices expose (usually intentionally minimal) being properly scrutinized with consideration for device management (by anyone, including the user)?

f) **Shared Responsibility.** How can shared responsibility and collaboration for IoT security be encouraged across stakeholders?

g) **Regulation.** Should device manufacturers be penalized for selling software or hardware with known or unknown security flaws? How might product liability and consumer protection laws be adapted or extended to cover any negative externalities related to the Internet of Things and would this operate in a cross-border environment? Would it be possible for regulation to keep pace and be effective in light of evolving IoT technology and evolving security threats? How should regulation be balanced against the needs of permission-less innovation, Internet freedom, and freedom of expression?

h) **Device Obsolescence.** What is the right approach to take with obsolete IoT devices as the Internet evolves and security threats change? Should IoT devices be required to have a built-in end-of-life expiration feature that disables them? Such a requirement could force older, non-interoperable devices out of service and replace them with more secure and interoperable devices in the future. Certainly, this would be very challenging in the open marketplace. What are the implications of automatic decommissioning IoT devices?

The breadth of these questions is representative of the wide-ranging security considerations associated with Internet of Things devices. However, it's important to remember that when a device is *on* the Internet, it is also *part of* the Internet,[61] which means that effective and appropriate security solutions can be achieved only if the participants involved with these devices apply a Collaborative Security approach.[62]

[61] Kolkman, Olaf. "Introducing Collaborative Security, Our Approach to Internet Security Issues." Web log post. Internet Society, 13 Apr. 2015. http://www.internetsociety.org/blog/public-policy/2015/04/introducing-collaborative-security-our-approach-internet-security-issues

[62] *Collaborative Security: An Approach to Tackling Internet Security Issues*. Internet Society, Apr. 2015. http://www.internetsociety.org/collaborativesecurity

Interoperability / Standards Issues

IoT Interoperability / Standards Background

In the traditional Internet, interoperability is the most basic core value; the first requirement of Internet connectivity is that “connected” systems be able to “talk the same language” of protocols and encodings.

Interoperability is so fundamental that the early workshops for Internet equipment vendors were called "Interops";[64] and it is the explicit goal of the entire Internet Standards apparatus centered on the Internet Engineering Task Force (IETF).[65]

Interoperability is also a cornerstone of the open Internet.[66] Barriers deliberately erected to obstruct the exchange of information can deny Internet users the ability to *connect*, *speak*, *share*, and *innovate,* which are four of ISOC's fundamental principles.[67] So-called "walled gardens'', in which users are permitted to interoperate with only a curated subset of sites and services, can substantially diminish the social, political, and economic benefits of access to the entire Internet.

In a fully interoperable environment, any IoT device would be able to connect to any other device or system and exchange information as desired. In practicality, interoperability is more complex. Interoperability among IoT devices and systems happens in varying degrees at different layers within the communications protocol stack between the devices. Furthermore, full interoperability across every aspect of a technical product is not always feasible, necessary, or desirable and, if artificially imposed (such as through government mandates), could provide disincentives for investment and innovation. The standardization and adoption of protocols that specify these communication details, including where it is optimal to have standards, are at the heart of the interoperability discussion for IoT.

Beyond the technical aspects, interoperability has significant influence on the potential economic impact of IoT. Well-functioning and well-defined device interoperability can encourage innovation and provide efficiencies for IoT device manufacturers, increasing the overall economic value of the market. Furthermore, the implementation of existing standards and development of new open standards where necessary help lower barriers to entry, facilitate new business models, and build economies of scale.[68]

A 2015 McKinsey Global Institute report states, "[on] average, interoperability is necessary to create 40 percent of the potential value that can be generated by the Internet of Things in various settings."[69] The report continues, "Interoperability is required to unlock more than $4 trillion per year in potential economic impact for IoT use in 2025, out of a total impact of $11.1 trillion across the nine settings that McKinsey analyzed."[70] While some companies perceive competitive advantages and economic incentives in building proprietary systems, overall economic opportunities may be constrained in a marketplace of silos.

[64] "A History of the Internet: 1988." Web log post. Computer Information, 12 Aug. 2010. Web. 6 Sept. 2015. http://inthistory4u.blogspot.com/2010/08/1988.html

[65] See http://www.ietf.org

[66] "Open Internet: What is it, and how to avoid mistaking it for something else," Internet Society 3 Sept. 2014. https://www.internetsociety.org/doc/open-internet-what-it-and-how-avoid-mistaking-it-something-else

[67] "Values and Principles." *Principles*. Internet Society, 2015. http://www.internetsociety.org/who-we-are/mission/values-and-principles

[68] The European Commission Rolling plan for ICT Standardisation 2015 section 3.5.6 Internet of Things has a discussion on IoT standards from a competitiveness and policy perspective. See https://ec.europa.eu/digital-agenda/en/rolling-plan-ict-standardisation

[69] Manyika, James, et. al., *The Internet of Things: Mapping the Value beyond the Hype*. McKinsey Global Institute, June 2015. p. 2. http://www.mckinsey.com/insights/business_technology/the_internet_of_things_the_value_of_digitizing_the_physical_world

[70] *Ibid.* 4.

Also, interoperability is fundamentally valuable from the perspective of both the individual consumer and organizational user of these devices. It facilitates the ability to choose devices with the best features at the best price and integrate them to make them work together. Purchasers may be hesitant to buy IoT products and services if there is integration inflexibility, high ownership complexity, concern over vendor lock-in, or fear of obsolescence due to changing standards.

Key Considerations and Challenges in IoT Interoperability / Standards

Interoperability, standards, protocols, and conventions are a primary issue in the early development and adoption of IoT devices. While not exhaustive, a number of key considerations and challenges include:

- **Proprietary Ecosystems and Consumer Choice.** Some device manufacturers see a market advantage to creating a proprietary ecosystem of compatible IoT products, sometimes called “walled gardens”, which limit interoperability to only those devices and components within the brand product line. These manufacturers can create user lock-in to their particular device ecosystem by increasing the switching cost for the consumer to change to a different brand in the future or substitute components from another vendor. For example, in the home automation market, light bulbs from one vendor may not be interoperable with a light switch or control system manufactured by another.

  Interoperability supporters view these practices as an impediment to user choice because it deters users from changing to alternative products. They also view this practice as a barrier to innovation and competition because it limits the ability of competitors to create new products based on the ecosystem’s foundational infrastructure. Some device manufacturers, however, see a closed ecosystem approach as a benefit to users by providing a protocol that can be adapted more quickly and easily when technical or market demands require change.

  Interoperability considerations also extend to the data collected and processed by IoT services. One of the primary attractions of connected devices is the ability to transmit and receive data to services “in the cloud”, which in turn provide valuable information or services based upon that data. While this is extremely useful, it also can present challenges for a user who wants to move to a competing service. Even if access to the data generated by devices is made available to users, obtaining the data will be useless if the data is in a proprietary format. Only if the source data is freely available to the originating user, in an open standard format, will users be free to move to another service provider, or to perform analyses on their own.

- **Technical and Cost Constraints.** As manufacturers develop IoT devices, there are inherent technical, time to market, and cost constraints that factor into device interoperability and design. Some devices are constrained by technical factors like limited internal processing resources, memory, or power consumption demands. Similarly, manufacturers are under pressure to reduce the unit cost of the device by minimizing part and product design costs. Manufacturers make cost-benefit analyses to decide whether the additional costs and potentially reduced product performance is worth the extra benefits of implementing standards. In the short-term, it can be more costly to design interoperability features into a product and test for compliance with a standards specification. In some contexts, the cheapest path to market may be to use proprietary protocols and systems.

This needs to be compared, however, against the long-term product lifecycle gains afforded by interoperability.

- **Schedule Risk.** In a globally competitive market, there is often a first-mover advantage to bringing a product to market quickly and establishing market share, and this certainly applies to IoT device manufacturers. A problem arises for IoT device interoperability when the device manufacturer's product design schedule outpaces the availability of interoperability standards. An IoT device manufacturer that is eager to bring a product to market may view lack of certainty in standards development schedules and processes as business risk to be minimized or avoided. This can make design alternatives to open interoperability standards more attractive, particularly in the short term.

- **Technical Risk.** When an IoT device manufacturer or user is planning the development of a product, they need to assess technical design risks of protocols in the development process. Incorporating existing and proven standards into product or system designs can represent a lower technical risk compared to the development and use of proprietary protocols. The use of generic, open and widely available standards (such as the Internet Protocol suite) as building blocks for devices and services can bring other benefits, such as access to larger pools of technical talent, developed software, and cheaper development costs. These factors are discussed in Internet Architecture Board (IAB) RFC 7452, "*Architectural Considerations in Smart Object Networking*". [71]

- **Devices Behaving Badly.** Lack of standards and documented best practices have a greater impact than just limiting the potential of IoT devices. In a passive way, absence of these standards can enable bad behavior by IoT devices. In other words, without standards to guide manufacturers, developers of these devices sometimes design products that operate in disruptive ways on the Internet without much regard to their impact. These devices are worse than simply not being interoperable. If poorly designed and configured, they may have negative consequences for the networking resources they connect to and the broader Internet.

  In an essay, Internet expert Geoff Huston describes the proliferation of such devices as the "Internet of stupid things''.[72] Huston describes an example of a consumer-grade cable modem produced by one manufacturer that hard-coded the IP address of the network time protocol (NTP) server operated by the University of Wisconsin into the product, which is a breach of commonly accepted design practices. As Huston explains, "The more units that were sold, the greater the aggregate traffic volume that was sent to the university's server."[73] Not only were these devices behaving badly by funneling all of the NTP requests to a single server, but the vendor's poor design compounded the difficulty because it provided no effective mechanism to fix the problem.

  There is an opportunity for the deployment of IoT standards and best practices to significantly diminish this class of problems over time

[71] Tschofenig, H., et. al., *Architectural Considerations in Smart Object Networking*. Tech. no. RFC 7452. Internet Architecture Board, Mar. 2015. Web. https://tools.ietf.org/html/rfc7452

[72] Huston, Geoff. "The Internet of Stupid Things." *APNIC Labs*., 28 Apr. 2015. https://labs.apnic.net/?p=620

[73] *Ibid*.

- **Legacy Systems.** Interoperability standardization is a challenge for new IoT devices that need to interface with systems already deployed and operating. This is relevant to many industry-specific and application-specific environments that have established networks of devices.[74] IoT engineers are faced with design trade-offs to maintain compatibility with legacy systems while still trying to achieve greater interoperability with other devices through the use of standards.

- **Configuration.** Users will face increasing challenges in managing larger numbers of IoT devices. One such challenge is the need to quickly and easily modify the configuration settings of many IoT devices on a network. When facing the daunting prospect of configuring hundreds of individual devices, it will be essential to have thoughtful design and standardization of configuration tools, methods, and interfaces.[75]

- **Proliferation of Standards Efforts**. Many new industry coalitions have emerged alongside traditional Standards Developing Organizations (SDOs) to increase efforts to assess, develop, modify, or harmonize standards and protocols related to IoT. This includes, for example, long-standing SDOs such as the IETF, ITU, and IEEE, and comparatively new efforts such as the Industrial Internet Consortium, Open Interconnection Consortium, ZigBee Alliance, and AllSeen Alliance, among many others.[76]

  The time and investment required by industry and other stakeholders to participate in the wide range of standardization efforts will likely be costly. Further, there is likely to be overlap and even conflicting standardization strategies between some efforts.[77] In addition to increasing the costs of standards development, the absence of coordination across efforts could ultimately produce conflicting protocols, delay product deployment, and lead to fragmentation across IoT products, services, and industry verticals.

Interoperability Questions

Interoperability and standards pose challenges and questions for the future of IoT devices, including:

a) In what areas are interoperability standards most needed and desirable? Are these sufficiently similar or different across the wide range of potential IoT applications and use cases (such as consumer goods, industrial applications and medical appliances)? What are the generic and widely available standards (such as the Internet Protocol suite) that could be used as building blocks for IoT devices and services? How would a lack of interoperability impact users' ability to connect, speak, share, and innovate?

[74] Examples of legacy system protocols include: SCADA (Supervisory Control and Data Acquisition), a protocol used for communication of industrial devices; CAN Bus (Control Area Network) protocols for vehicle and industrial sensors.

[75] Vint Cerf, personal communications, 9 September 2015.

[76] See section "For More Information" at the end of this paper for a list of standards bodies, consortiums, and alliances working on IoT standards issues.

[77] Lawson, Stephen. "Why Internet of Things 'Standards' Got More Confusing in 2014." *PCWorld*, December 24, 2014. http://www.pcworld.com/article/2863572/iot-groups-are-like-an-orchestra-tuning-up-the-music-starts-in-2016.html

b) What are the optimal roles of Standards Developing Organizations (SDOs), industry consortia, and stakeholder groups in IoT standards development? What is the potential for bringing together the wide range of groups working on IoT technical implementations for broader discussions about interoperability and standards implementation? Can competing standards, duplication, and conflicts stemming from SDOs and consortia tackling similar or overlapping issues be avoided without adding undue coordination overhead? More practically, how can industry players and other interested parties keep track of all of the activities happening in this broad space?

c) What is the best approach to educate and engage user and developer communities about the problems of badly behaving IoT devices and lack of standards implementation? What types of best practices or implementation reference models would be effective, given the broad range of IoT applications and use cases?

d) How will the Internet of Things impact the consumption of bandwidth and other resources and to what extent will standards need to be modified to support those evolving needs? Given the importance of cloud-enabled services to the Internet of Things, what are the challenges related to cloud-to-cloud interoperability?

Overall, the importance of IoT interoperability and standards to the market and consumers is undeniable. Ultimately, the challenge of developing and employing interoperability standards is central to the discussion of innovation, competition, and user choice of services, which are embedded in ISOC’s core principles.