Internet of "What?"

G. Hurlburt (Change –Index, Inc.)
J. Voas (US National Institute of Standards and Technology)
K. Miller (U. of Illinois-Springfield)

Things.

The Internet began as a communication medium between a fairly restricted set of people. The development of the web turned the Internet into a communication medium between people and organizations, and soon after that, between organizations and organizations. Today, we see a new Internet player becoming more and more important: things. A "thing" in this setting is any inanimate object that can be programmed to communicate, sense, and interact with other things.

So what can Internet things be? Home appliances, any type of sensing device, an automobile, a "smart flashlight", and even "smart doorknobs"[1] have all become candidate " Internet things."

There are now far more "things" connected to the Internet than there are people who use the Internet, and that is only counting PCs, Laptop, tablets and smartphones. Internet things have outnumbered Internet people since 2008[2]. As IPV6 comes into common use, the available Internet Protocol (IP) address space becomes a very big number, 2^{128}. This number well exceeds the 7 billion people presently occupying our planet. Clearly, "things" can occupy most of these IP addresses - smart "things" that can always be traced back to human users, but sometimes by a very indirect route. Moreover, the range of what these things do is huge and growing by the day.

The RFID community envisions total visibility into the supply chain to permit unprecedented real-time distribution flexibility. There are already success stories in this field, including extensive use by FedEx .[3] The European Union (EU) foresees real-time health monitoring for all using RFID.[4]

Along with the EU, many large corporations predict all manner of real-time sensors on the Web for utility management, traffic control and other urban efficiencies. Others see significant advantages to households as Internet things proliferate, even just from an energy management standpoint. IBM is actively promoting an Internet of Things (IoT) protocol within the EU.[5] Cisco speculates on the broader Web of Things (WoT), thinking of the interactions between autonomous electronic devices. And, of course all become the recipients of the goodness wrought by networks of these inter-connected, interactive webbots operating in clouds on our behalf.

The literature promoting the IoT frequently uses the example of a hypothetical harried businessperson who must make an early morning meeting on time. The objective is for this individual's Internet savvy alarm clock to go off in time to allow our pre-occupied executive to make a specific meeting - well rested and without hassle. In such scenarios, the faithful clothes dryer will have completed its cycle just in time to provide the appropriate attire for the day. The toaster and coffee maker will have smartly collaborated to produce a warm breakfast snack and presumably, our commuter's car's passenger compartment will be air conditioned for the morning run to the train station. Of course, interconnected highway sensors have successfully predicted the commute time to the train station. Equally coupled railroad systems pinpoint the arrival time of the train suitable to get our commuter to the meeting. All of these events and, many more, feed information to the smart alarm clock such that it can self-synchronize to trigger actions around the house and ultimately "go off" at just the right moment.

Unfortunately, this rosy scenario may not accurately predict the reality of relying on the IoT. Perhaps when they emerge from the drier, the clothes sorely need ironing. Our commuter may have neglected to pre-load bread into the toaster the night before. A coffee maker relay fails, and the coffee crusted meltdown is taking shape inside the smoky coffee pot. A phone call from a far less well-connected colleague who is mired in traffic behind a fresh major accident costs valuable time. The car needs gas, but the alarm clock could not make such an inquiry from the older model car, which had no intelligence about its dire fuel state. After a ten-minute detour to fill up makes our slightly disheveled, hungry, caffeine-starved commuter late for the last train that could possibly have delivered him to the meeting in time. Needless to say, vendors of IOT equipment will insist that the answer is clearly more sensors everywhere, more smart machines and better coordination of the things to get our commuter effortlessly to work. This, in turn, makes the promoters of IOT exceedingly happy. But will an increasingly fragile ecosystem be able to sustain the amount of power necessary to run all these gadgets? And so it goes.

The foundation of the IoT is data. The more that intelligence sensors, RFID tags and smart devices are attached to the Internet with their own IP addresses, the more data there will be to be handled. The amount of data will continue to grow exponentially, furthering the belief that we are increasingly swamped with raw data. The underlying question then becomes one of harnessing all of these data into something intelligible. Networking certainly helps in routing raw data for subsequent interpretation and action, often under autonomous control within a well-defined domain. For example, we have already seen that road sensors can paint traffic patterns that can drive traffic lights to smooth the flow of vehicles based on real-time events. Further filtering and routing assists in getting some semblance of information to decision makers, sometimes human, but sometimes not human. As another example, the actual flow of traffic overnight, as melded into patterns established over time, serves to predict the next morning's flow. Assuming royalties are somehow paid, this predictive information, in the form of yet more data, will make their way to our commuter's well-connected alarm clock to help

determine the optimal wake up time. At some point, however, human behavior will be affected either positively or negatively. That is the very real and rather scary promise of IoT. “We shape our tools and thereafter, they shape us”.[6]

Approaching the classic Gartner “Peak of Inflated Expectation”, IoT hype is just beginning to have an impact. [7] (See Figure 1.) Interestingly, IoT is appropriately preceded by Private Cloud, Cloud computing and Machine-to-Machine (M2M) communications, all of which are sliding towards Gartner’s “Trough of Disillusionment” before finding their rightful places in the broader scheme of things automated. If Gartner’s curves are credible, the hype for IOT will ratchet up to a crescendo before practical implementation gains any real traction. This suggests the luxury of some time, perhaps 5-10 years, to deal with IOT rationally. Some real issues exist.[8]

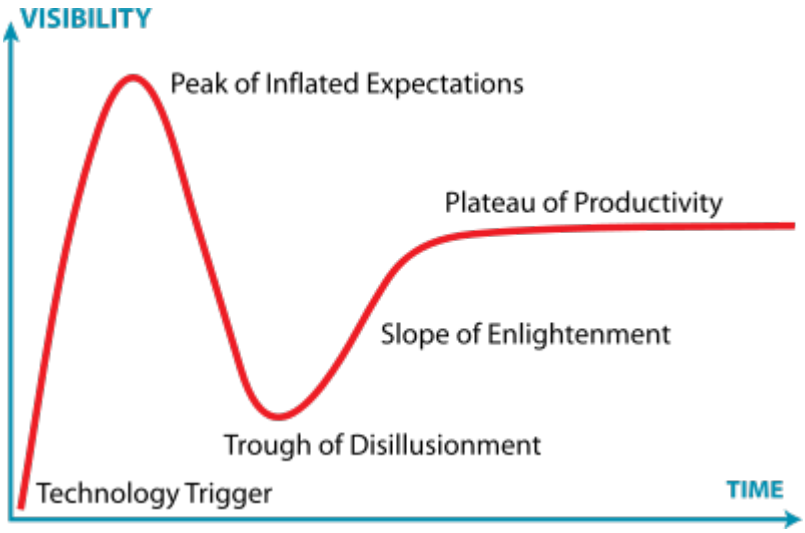

Figure 1. Gartner’s Hype Cycle [1]

The sheer amount of data generated by an IoT eventually becomes unfathomable. As noted, how these data are routed and managed will help ease the questions of their utility. British tabloids allegedly hacked private texting, a large body of data. Any IoT data flowing over unprotected public networks are vulnerable to this kind of thievery at the expense of those exploited. Thus, the cyber-security stakes for IoT data in transit spiral upwards.

Moreover, some serious questions arise when we consider how these data are converted and harnessed for useful purposes. Absent standardized protocols, how would a smart alarm clock, for example, know to program all the peripheral household devices necessary to maximize the odds of just in time delivery of our executive to a meeting? Somehow, the clock must be aware of the commitment time for arrival suggesting some sort of remote schedule coordination and location awareness. It must be in minute-to-minute contact with household appliances to coordinate their activities. It must rely on data likely selectively pulled from outside the home network, to make

critical trip planning assumptions. Lacking an overarching data protocol stack, such coordination would be piecemeal at best. While IPv6 will manage the transport layer, this problem reaches to the heart of the "Application Layer" atop the famous ISO hierarchy of network layers.[9] This kind of close coordination has proven elusive in practice today; the goal of universal implementation in most domains seems unlikely to happen in the foreseeable future.

As the IoT becomes ubiquitous, issues of information ownership will become crucial. Who will own the oceans of data IoT will generate? There are already monumental battles over who owns medical data and under what conditions may they be shared, and this is well before they become ubiquitous. Interestingly, medical data are not ubiquitous because the various lexicons, tight as they may be, are not yet sufficiently standardized to allow sharing without risk of serious misinterpretation. This suggests that ownership and lexical interpretation of data have an economic death grip on one another that may not be easily broken soon. Even if this thorny problem is resolved, another firestorm is brewing over data transfer between public, hybrid and private clouds, a problem that will become even more acute as the IoT expands.

Even if data are structured in such a way that enables clean transference, there is a cost for doing so. Consider the case of instrumenting both municipal roads and private throughways to monitor and control traffic flow. Such a system, entailing an extensive sensor-based infrastructure and a highly sophisticated processing commitment, requires resources to create and maintain. It is highly doubtful that such data would pass the portals of anyone's smart alarm clock without some form of compensation, be it in a service fee or added taxes. The means to meter data as a function of value received is still very much in its infancy; the kinds of conflicts that will arise can be seen in the tumultuous publishing and music industries as the Internet changes fundamental assumptions about who owns what. Moreover, such issues beg the questions of ownership and copyright, as IoT data may well have multiple uses hardly conceived upon their initial provisioning. Another complication will arise as data from different sources is accumulated, analyzed, and then sold. How will each data source and each analyzer be compensated for the appropriate amount of value added? If data sources and analyzers do not think they are receiving fair gain for their efforts, it seems unlikely that the data will continue to be collected and analyzed.

Existing skirmishes between privacy and security concerns may blossom into a battle as the IoT expands. The British and other Europeans seem to have accepted total public surveillance as a way of life. Surely, with facial recognition, localization and commercial transaction capabilities rapidly gaining velocity, it will soon be nearly impossible for any citizen to live off the grid. Most transactions and activities are being captured digitally, recorded and utilized to lubricate commerce and to facilitate tighter security controls. For the convenience of doing business, much of everyone's identity is transferred into digital form, and we all run the risk of having the details of our lives become a commodity that is bought and sold without any individual control or gain. Checks and

credit card transactions are being digitized and tracked and reconciled at or very close to the time of the transaction. Is there a point at which harvesting identity for commercial gain will become intolerable? And it is not only commercial forces who are hungry for our digital data souls; governments are also pushing for more and more information about individuals, fueled by and masked by a post-9/11 mania for protection against enemies who are often vaguely defined, non-traditional combatants.

As machines begin to reason about our individual personae, legal constraints against unreasonable electronic search and seizure need to be shaped and enacted. A balance must be struck between the right to individual privacy and the security of the crowd to assure that moment-to-moment reasoning, involving mountains of sensor and transaction data, cannot exceed reasonably defined legal thresholds without risk of penalty.

These issues, heady as they are, become almost insignificant in the larger IoT picture. By definition, an IoT is a network phenomenon. Early network research involving the Internet established that a few hubs are critical to connections between nodes. Think of the role of Facebook, Google, Amazon and Groupon in forming connections throughout the Internet, much less throughout society. Recent studies tell us we are each distant from everyone else on "Facebook" by 4.74 degrees of separation, as opposed to the initially projected estimates of around 13.[10] Moreover, the few most important hubs gain in influence as the network expands, making them critical assets to network survival. This centrality of the most important hubs is supported both by empirical data, and by theoretical explanations. While this exposes a vulnerability of sorts, there is a far deeper implication, just coming to light.

Recent research now couples insights about the Internet with sophisticate control theories. This suggests that network characteristics can be intentionally manipulated for a purpose.[11] While morally agnostic, this is a powerful concept, as the purpose can be for good or ill in the eyes of the beholder. This is somewhat mitigated by the ongoing private, public, hybrid cloud model debate, but its portent serves as a wakeup call to solve the foregoing problems. We must be vigilant lest the IoT leads us into an electronic future that harnesses our resources into what amounts to an Orwellian nightmare.

REFERENCES

[1] V. Shannon, "Wireless: Creating Internet of 'Things': A scary, but exciting", New York Times, Technology Section, Nov 20, 2005.

[2] D.Evans, "The Internet of Things - How the Next Evolution of the Internet Is Changing Everything", Cisco Internet Business Solutions Group (IBSG), April 2011.

[3] Editorial Staff, "RFID News: Somewhat Behind the Scenes, the "Internet of Things" is

Moving Forward", Supply Chain Digest, Dec. 09, 2009

[4] Commission of the European Communities, "Internet of Things — An action plan for Europe", Communication From The Commission To The European Parliament, The Council, The European Economic And Social Committee And The Committee Of The Regions, July 18, 2009

[5] C. Rubsamen and H. Tomasson, "Eurotech and IBM Contribute Software to Connect Next Generation of Wireless and Mobile Devices, Eclipse Contribution to Create New Standard that Connects Internet of Things", MarketWatch, Nov. 3, 2011

[6] M. McLuhan, "*Understanding Media: The Extensions of Man",* Gingko Press, 2003

[7] Gartner, Inc. Hype Cycle Research Methodology. http://www.gartner.com/technology/research/methodologies/hype-cycle.jsp, accessed 6 Dec. 2011.

[8] j. Brockmeier, "Gartner Adds Big Data, Gamification, and Internet of Things to Its Hype Cycle", ReadWrite Enterprise, Trend Analysis, Aug. 11, 2011

[9] H. Zimmermann, "OS1 Reference Model-The IS0 Model of Architecture for Open Systems Interconnection", IEEE Transactions .On CommunicationS, VOL. COM-28, No. 4, APRIL 1980

[10] J. Markoff and S. Sengupta, "Separating You and Me? 4.74 Degrees", New York Times, Technology Section, Nov. 21, 2011

[11] Y. Liu,J. Slotine3,4 and A. Barabasi, "Controllability of complex networks", Nature, Vol 473, p 173, May 12, 2011

Disclaimer: This paper was co-authored by Voas as a private citizen and not as a US National Institute of Standards and Technology (NIST) employee. No NIST resources were used. It reflects Voas's personal opinion and does not reflect the opinions of the Department of Commerce or NIST.