EVOLVING TO AN ECOSYSTEM OF IoT APPLICATIONS

PILGRIM BEART

Evolving to an ecosystem of IoT applications.

Today's vertical Machine-to-Machine silos are a major inhibitor to the arrival of the true 'horizontally connected' Internet of Things, creating vendor lock-in and inhibiting choice, says Pilgrim Beart, CEO of 1248 Ltd. Here he examines possible solutions and analyses the fundamental trends driving the advent of the Internet of Things along – with both positive and negative consequences.

Because many IoT applications have been designed to reduce resource utilisation, the growth of the IoT ecosystem could be to everyone's common advantage, enabling a win-win-win benefit for customers, suppliers and government agendas alike.

What is IoT really?

The Internet of Things, is a term which has recently gained huge popular currency. However, so broadly has it been applied that it is becoming devalued as a term. So before diving in, it is worth considering how – if at all – IoT might really be different from anything that has come before.

The potential for IoT would seem to be primarily to do with scale. The number of connected devices on the planet is growing – at about 25% per year according to Cisco [1]. Given that the population is not growing at anything like that rate, the compound effect of that growth is that the number of connected devices per person on the planet is growing about tenfold per decade.

This can be seen in our personal lives: a decade ago we had just one connected device – our PC. Today we have around ten connected devices – laptop, smart phone, tablet, connected car, Nest or Hive thermostat, Sonos sound system, Sky+ or AppleTV box and so on. In another decade, we'll have a hundred connected devices, and a thousand a decade after that. That same trend is happening outside the consumer sphere too: in industry, commercial and public ecosystems.

Admitedly, this adoption is driven partly by "technology push", such as falling prices and easier connectivity. But it is also being driven by "customer pull", as users seek benefits such as convenience, control and automation.

Exciting though the potential is, like any fast-moving technology, it is not without its challenges. As any consumer will know, today's connected devices are driving a proliferation of incompatible gateways, services and apps, threatening to make everyone's lives more complex. And the growing number of devices is driving a growing need to manage them through their lifecycle – something the end-user may not have the time, knowledge or inclination to do.

These devices will need to manage themselves in a way that our creations have never had to before.

Applications of IoT

It seems likely that IoT will become as pervasive and transformational over the next 15 years as the Web has over the past 15 years. The answer to the question "What are the applications for the IoT?" is probably "everything – and more that we've never even thought of". However, by focusing on particular ecosystems, we can explore how IoT might emerge from today's siloed connected-device applications.

One such ecosystem is the HyperCat[1] consortium, a UK initiative supported by the Government, which is in the process of developing twelve so-called "spearhead" applications, all with a Smart City/Smart Society theme. These are broadly divided into four functional segments, as shown in Figure 1 overleaf, all of which are potentially huge global markets for IoT.

Manufacturing segment – determining where things are (asset management), tracking production to enforce quality, and food safety by tracking food temperature through the supply chain. The general focus is on reducing costs while increasing efficiency, quality and compliance.

Transportation segment – fleet management, including rapid diagnosis and reporting of vehicle faults, supporting drivers with relevant data during the journey (mashing-up live data from many sources including the Highways Agency, Weather and Vehicle behaviour to help drivers drive more safely and economically).

Utilities/Energy segment – measuring and managing large users of energy, both direct and indirect. Connected streetlights can be managed to optimise both consumption and repair schedules (with cost savings estimated to be circa £4m for a city the size of Milton Keynes). Electricity storage at grid-edge, in batteries, enables householders to increase their self-use of solar electricity by about 25% providing benefits to the Distribution Network Operators who run the last mile of distribution. Pumping water is a big indirect user of energy too, and live tracking and management of water levels in rivers can save energy as well as avoiding flooding.

Smart Cities segment – tracking and managing assets in buildings, including the various parts of Building Management Systems. It also includes Smart Parking (sensors in parking bays, apps to recommend free spaces) and Smart Mobility through an application called My Guardian (safety for travellers, for example ensuring that students travelling to and from campus reach their destination safely, and raising the alert quickly if not).

As is common in IoT, each of these applications turns dumb products into connected products. Moreover, it enables (and increasingly mandates) a shift in business model, from a product model (up-front capital expenditure) to a service model (ongoing operation expenditure).

A closer look at one application – Utilities/Energy

All these sectors are undergoing rapid change, driven by changes in technology and by macroeconomic shifts such as socio-demographic shifts and global change. However, examining one sector in more detail illuminates the potential use of IoT and some of the challenges.

If IoT is about innovation – and it is – then the Energy sector is a particularly fertile place to make use of it. The whole of our energy infrastructure is being reinvented in a once-a-century upheaval to address climate change, including a shift from fossil fuels (easily stored) to electricity (not), and a shift from demand-led generation (turning up a fossil-fuelled power station quickly to respond to a demand peak) to a world of supply-led demand (if the wind isn't blowing and the sun isn't shining, then consumers need to respond to this lack of supply).

This is necessitating change up and down the supply chain in what has historically been a rather conservative industry, changing the approach to generation, national distribution, local distribution and end consumption. The old centralised command-and-control paradigm, based largely on forward planning and sending faxes around, breaks down as generation becomes distributed and supply becomes increasingly non "dispatchable". Politicians may talk of an energy trilemma (delivering secure energy supply at reasonable cost whilst decarbonising) but the word 'trilemma' implies a choice; in reality there is no choice – all three must be achieved.

One company working on an interesting solution is Moixa Technology[2], which has

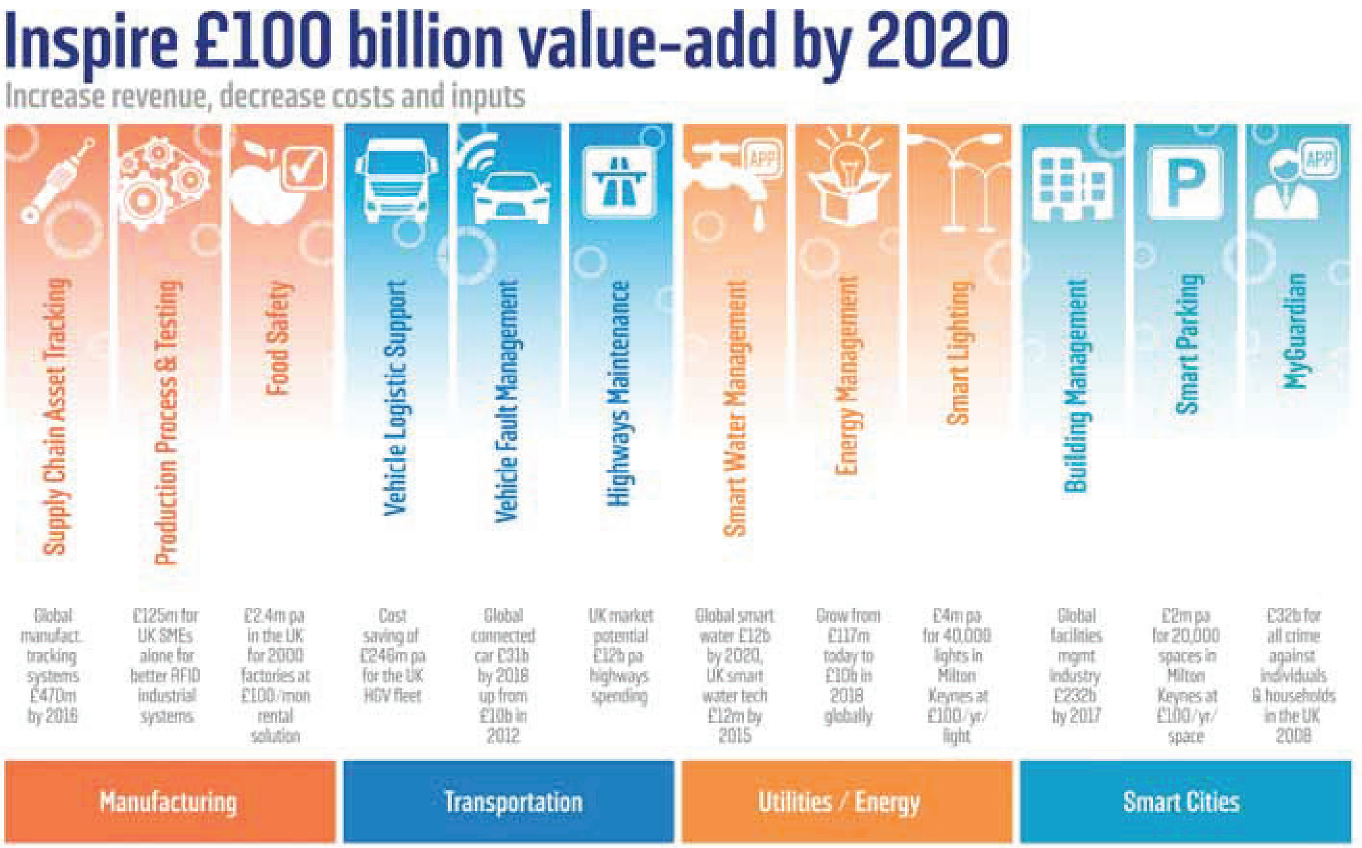

Figure 1: The "spearhead" applications of the HyperCat Consortium

developed a grid-edge electricity storage solution called Maslow – a 2kWh of managed battery storage, with an inverter, packaged into a box which looks something like a small domestic wall-fitted boiler. It is designed primarily for homes with solar photovoltaic panels. Such homes typically only consume about 25% of the electricity they generate, because the sun often shines when there is no-one home using the electricity, and then often isn't shining during the evening peak. Maslow addresses this by storing excess day-time consumption, giving it back at the evening peak, increasing self-use of solar electricity from about 25% to about 50%, i.e. doubling the value (in pure energy terms) of the panels to the home-owner.

An interesting aspect of Maslow, shared with many IoT solutions of the author's acquaintance, is that it has potential benefits for all the many players in the energy supply chain. Moixa estimates the value to the home-owner is sufficient to give it a nine-year payback (even without subsidy, and before time-of-use pricing arrives). But it also has value to the local distribution operator, responsible for the electricity cables in the street, because a few well-placed Maslow systems can prevent those cables going out of voltage tolerance when there is a high level of feed-in from photovoltaic panels, or a lot of drain from heat pumps or recharging electric vehicles. The investment in digging up the road and replacing substations might better be spent on giving a few homeowners a Maslow. And for the national network, the ability to control this storage, so it can store excess electricity e.g. from wind farms, and help address generation outages, is of real potential benefit too (the National Grid pays around £60,000 for every MW of load that can be made dispatchable). So there is a real win-win for all players, including the Government in addressing its so-called energy trilemma.

All this of course only relies on connectivity. Each Maslow is "smart" and connected to the internet. This is necessary so that:

- Its effectiveness can be monitored by stakeholders including homeowners, housing associations, the manufacturers, and even grid operators.
- It can be controlled remotely, for example to change how it is programmed, or to dynamically cause it to fill or empty its battery store.
- It can be managed and maintained remotely (particularly important in trial situations where vendors are learning how devices are used).
- It can report its status, its readiness to dispatch load, and report-back measurements of how much power it actually delivered in any period.

Connectivity with the Maslow can be via WiFi or mobile cellular which other applications can then exploit – for example Moixa are experimenting with pooling community energy use, allowing neighbours to share spare solar electricity – implemented as an accounting practice by sharing data.

Energy is a market where there is a close relationship between "device functionality" and "getting paid" and so the presence of a live connection to demonstrate that each unit is functional is vital.

As silo walls come tumbling down, islands grow and connect

Each of the above applications will benefit if vendor silos are broken down in favour of horizontally integrated systems. Such systems bring benefits to the users (more choice of provider, less vendor lock-in, and more competitive pricing) and also to the vendors (larger, more homogenous markets leading to economies of scale).

As these islands of interoperability become established and grow, opportunities emerge to connect one island to another; the move to the IoT becomes an avalanche. It is not hard to see how connecting data between any of these sectors can be of mutual benefit – to make up one example: streetlights could light only bays in which cars are parked.

A common theme is that the IoT has enabled better use of resources, whether they are parking spaces, energy, or peoples' time. This in turn drives opportunities for mutual benefit amongst all participants (typically the public, government and vendors), potentially creating a win-win-win dynamic which can pull the applications and inter-connectivity through to widespread adoption. That's the theory, anyway.

Why isn't IoT here yet?

There is no fundamental obstacle to IoT, no rocket science needed. The potential benefits seem substantial. But there are at least a couple of thorny problems:

Standards – One is a lack of interoperability driven by the multiplicity of technical standards at every level of the stack. There is no lack of standards – quite the contrary. As Andrew Tanenbaum once said, "The nice thing about standards is that you have so many to choose from" [2]. In physical wireless connections alone there are more than four hundred different standards and this multiplicity of standards makes it statistically unlikely that any two connected device applications will "just happen" to be compatible.

In an attempt to address this, the previous decade has seen some top-to-bottom standards, such as ZigBee[3], which specify every part of the communications stack, from physical link to application interface, in an attempt to nail everything down and thus create a homogenous, interoperable world. However, these have typically suffered from the problem that, in order to realise their benefit, you HAVE to use the whole standard, and they are very large and complicated, thus hard to implement and hard to evolve to ensure they remain the best of everything.

Naming – Another problem at the higher levels of the stack is that of "ontology" – in order for one computer to exchange data with another, the two have to agree on common naming. For example, if an application wants to make sense of data from a temperature sensor, then both need to share a common understanding of, and literal name for, the concept of temperature. If one calls it "temp" (in Fahrenheit) and the other "degrees" (in Centigrade) then, even if they can exchange data, it will not be meaningful.

How to accelerate IoT?

First and foremost, it is the INTERNET of Things. All the interoperability, facilitated by the Internet Protocol (IP) which already pervades everyone's pockets, desks and buildings, should be leveraged to the maximum.

The extreme version of this vision is IP-to-the-edge – an interesting trend in modern standards such as Thread[4]. It promises the vision – already realised with the Cellular and WiFi IP connections on our smart phones – that it matters not to apps or users how a connection is made, as long as it's IP. Though a worthy goal, there are still various practical obstacles to truly bringing IP to the edge, not least that IP's standard connection-oriented protocol (i.e. it's Transmission Control Protocol) cannot cope with sleepy IoT end devices or the latency of narrow-band protocols such as Weightless[5]. However, provided there is transparent and consistent translation between IP and nearly-IP-at-the-edge standards, then many of the benefits of IP can be realised without having to send IP packets right to the edge.

Another principle that can be taken from IP is that it approaches problems in a layered fashion. New standards in the world of IP attempt only to fill existing gaps, rather than inventing a whole new stack. IP standards are also often made in a de-facto, after-the-fact fashion through the Internet Engineering task Force / World Wide Web Consortium, which is rapid and agile. Contrast this with the telco-style up-front standards-making approach which successfully brought us, for example, the various generation of mobile cellular standards, but is terribly slow.

Simplicity

From the perspective that matters most – that of the end-users – the biggest challenge of all for IoT is perhaps usability – making IoT "just work" (and stay working!). As IoT devices become more numerous and more diverse, it becomes more and more unreasonable to expect users to understand how to look after them, or even to have the time to look after them. It's hard to make things simple, and this is true for Things too. The field of User Experience, once viewed as an optional nice-to-have niche specialism of relevance only to a few, takes centre-stage with IoT, because Things will simply not be usable unless they have been made simple [3]

Interesting initiatives

As one might expect, there are many interesting initiatives currently addressing some of the above challenges. Some are addressing specific problems, usually with work that will in time become technical standards thereby sweeping away interoperability obstacles. Others are large consortia, aiming to identify and solve common problems, largely by selecting existing standards, driving IoT's arrival by creating critical mass within an ecosystem. These initiatives are sometimes portrayed in the media as competing and by implication incompatible whereas in fact there is a large degree of commonality between the approaches and standards being chosen, leading to some optimism that the IoT may indeed be just around the corner.

Three current initiatives are currently of particular interest to the author, all being conducted in an open, agile fashion.

IP-to-the-edge – In the lower layers of the communications stack, the IP-to-the-edge trend continues, driven largely by initiatives such as Thread, a mesh wireless Local Area Network building on the same IEEE 802.15.4 [4] physical layer already proven by ZigBee, but adding IP support by using IPv6 over Low power Wireless Personal Area Networks (6LoWPAN) specified in RFC 4944 [5], plus standard management layers such as defined in RFC 7252 [6] and the Lightweight Machine-to-Machine device management protocol defined by Open Mobile Alliance [7].

Thread is a completely new standard, driven by consumer giants such as Google Nest[6] and ARM[7]. Meanwhile other standards such as Bluetooth are becoming IP-enabled. The interesting possibility here is that, in the consumer domain, individual companies such as Google or Apple have sufficient might to drive millions of units out in the world unilaterally, thus short-cutting the normal slow ecosystem boot-up process. This may have significant benefits even to applications beyond the consumer space, where these standards and chip-sets can be adopted once proven at scale.

Beacons – Apple launched iBeacon[8] as a very simple local Bluetooth-based beaconing system – each beacon just emits a short unique code, which a smart phone can then look-up in order to launch a Web page. Google is now evolving this thinking into what they call the "Physical Web" by allowing beacons to emit URLs directly, and providing a mechanism for Web pages to send messages back to beacons using Javascript embedded within the page. Example use-cases include paying for ones' parking meter via the Web. Thus the smart phone becomes a kind of mobile gateway or "mother ship". Google seem to be aiming towards a world where the Web RESTful API[9] is the common interface for IoT.

HyperCat – This is the consortium used above as an example of an application ecosystem. Technically, it is focussing on the particular problem that there are an exponentially-growing number of IoT clients, and an exponentially-growing number of IoT services, and users increasingly wish to apply clients to servers in any combination. This combinatorial explosion is hampered by the fact that – even though clients and servers generally all support common interface standards such as Hypertext Transfer Protocol (HTTP), JavaScript Object Notation (JSON) and RESTful APIs – a human must still write some software to enable every combination. Humans are therefore the bottleneck and HyperCat seeks to address this by providing an automated way for IoT clients to discover data they can understand in any IoT service.

AUTHOR'S CONCLUSIONS

The trends driving the IoT appear to make it inevitable, bringing both increasing power for good, but increasing risk of complexity to address. Vertical silos and technical fragmentation are currently inhibiting IoT, and this is being addressed through a combination of technical standards-making and proactive ecosystem-creation.

ABOUT THE AUTHOR

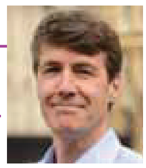

Pilgrim Beart FIET

Pilgrim is a computer engineer turned technology entrepreneur. He previously co-founded several "connected product" start-ups - activeRF, antenova, Splashpower and AlertMe, the connected home platform which manages millions of devices and powers British Gas' Hive, sold in 2015 for $100m. Pilgrim is now CEO of 1248, whose flagship Device Pilot helps connected-product companies to manage their products at scale.

ABBREVIATIONS

6LoWPAN	IPv6 over Low power Wireless Personal Area Networks
API	Application Programming Interface
IoT	Internet of Things
IP	Internet Protocol

REFERENCES

1. Cisco. How Many Internet Connections are in the World? Right. Now. Jul 2013.
2. Tanenbaum, A. Computer Networks, second edition, p254. Sep 1988
3. Rowland, C., et al. Designing Connected Products: UX for the consumer internet of things. O'Reilly, Apr 2015
4. IEEE 802.15.4. IEEE Standard for Local and metropolitan area networks - Part 15.4: Low-Rate Wireless Personal Area Networks.
5. IETF RFC4944. Transmission of IPv6 Packets over IEEE 802.15.4 Networks. Sep 2007.
6. IETF RFC7252. The Constrained Application Protocol (CoAP). Jun 2014.
7. Open Mobile Alliance. LightweightM2M. Oct 2015.

FOOTNOTES

1 See: http://www.hypercat.io/
2 Moixa Technology. See: http://www.moixatechnology.com/
3 ZigBee. See: http://www.zigbee.org/
4 Thread. See: http://threadgroup.org/
5 Weightless. See: http://www.weightless.org/
6 Google Nest. See: https://nest.com/uk/
7 ARM. See: https://www.arm.com/
8 iBeacon. See: http://www.ibeacon.com/
9 RESTful API. See: https://en.wikipedia.org/wiki/Representational_state_transfer

ITP INSIGHT CALL

Want to talk to the author?

To discuss this article and its content, join in the ITP Insight Call on 8 January, 2016.

To book onto the call visit: https://www.theitp.org/calendar/