CREATING **INTERNET OF THINGS** ECOSYSTEMS

MIKE FISHER, JOHN DAVIES

IoT ecosystems, a case in point.

Just what constitutes an Internet of Things (IoT) ecosystem and how does it work in practice? Mike Fisher and John Davies examine some of the practical issues associated with creating an IoT ecosystem, using the city of Milton Keynes as their example.

The Internet of Things (IoT) is anticipated to deliver significant innovation in many different areas, including future cities, transport, health and social care, manufacturing and agriculture. Sensors can be deployed at low cost to measure and monitor to a far greater extent than has been possible before. There is increasing recognition of the potential value in opening up data resources so that they can be exploited more fully. The ability to achieve more effective communications with people wherever they are adds an extra dimension to the potential applications that can be built. There is increasing activity and interest in exploring this rich set of possibilities around the world. However, the IoT exhibits considerable diversity and complexity which means that significant barriers to innovation remain.

At the highest level, many of the IoT applications being considered appear similar – involving the collection of information from a range of sensors and other sources, interpreting this in a specific context, and then making better decisions which improve behaviour compared with what was done before. There are two important classes of sensor-based IoT application – those which aim to monitor and respond to time-sensitive conditions, and those which collect data over a period of time for analysis. In either case, much of the time and effort involved can be spent on activities which appear generic. It would be much easier to build a new application in an environment where these generic problems have already been solved by others, with robust solutions available to all. This is what we mean by an IoT ecosystem.

IoT ecosystems

From a technology point of view, there is little fundamentally new in IoT – it is all about working at a different scale (more devices, more data, more scope for automation), and the potential to share information much more widely, driven by ever-decreasing component cost and device miniaturisation. There are a number of important new challenges, however. The potential advantages of information sharing are accompanied by real concerns over security and privacy. An IoT ecosystem consists of a number of independent stakeholders sharing a common interest in particular kinds of information and obtaining overall benefit from participation. This could be as a commercial provider of information or analytic services, as an application developer or end user, for example.

Use of shared services and facilities generally involves a compromise – typically giving up some level of direct control in return for reduced costs. In the case of today's global communications networks (including the Internet), the case for common services is very strong. Cloud computing and storage are also becoming widely accepted, although there are still many situations where private infrastructure is preferred. The potential for an IoT ecosystem to stimulate and enable
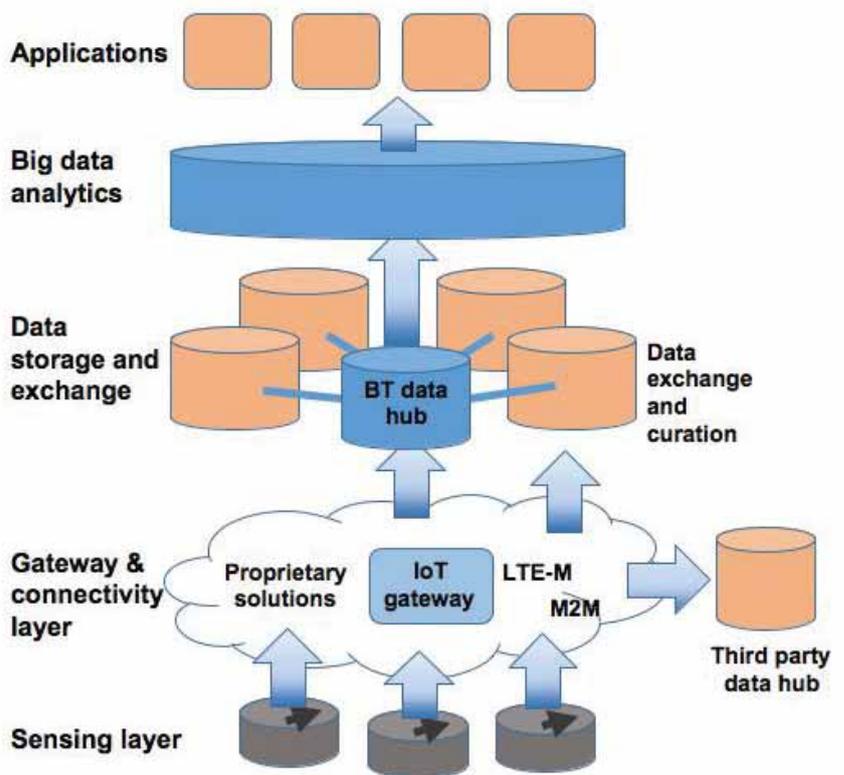

Figure1: IoT ecosystem value chain

innovation is clear, but all participants need to have confidence in the value proposition and be convinced that it meets their needs. If this is not the case then a sustainable ecosystem will not be possible. Areas of concern will include security and trust, respect for personal and commercial rights, dependability, performance, the ability to comply with legal and regulatory obligations, and cost. Predictability, simplicity and flexibility are particularly important characteristics.

Participants in an IoT ecosystem

Stakeholders in an IoT ecosystem may take one or more roles. These include information providers, application developers, analytics service providers, platform providers, and users of information and applications. Information providers and users can be either individual or organisational. These roles are shown in the value chain of Figure 1 and are described in more detail below.

Information providers

Information providers in IoT are often owners of sensor deployments. The primary purpose

of their sensors may be for their own use but they may choose to make some of their data available to others, either on a commercial basis, to meet their obligations (particularly for public sector organisations), or for the general good, or on a commercial basis. Other types of information are also of value in IoT applications, even if not directly associated with "Things". These include contextual (e.g. geographical, administrative) information and notifications of events such as traffic incidents, sporting fixtures etc, both of which can help make sense of sensor data.

Information providers should be able to publish data resources and advertise their availability via an easily accessible catalogue so that potential users can independently discover and assess their utility. Making data available should not imply relinquishing ownership rights so the information provider needs the ability to define access controls as well as terms and conditions for use of their data.

A new role is anticipated as the ecosystem develops – that of the derived information

provider which, while not necessarily being the primary source of any information, adds value by combining data from multiple sources, transforming or applying various analytical techniques. The derived information can then be republished.

Application developers

Application developers produce applications which make sense of available information within a specific context, and which provide value to end users. They should be able to discover what data resources are available to them, and assess which ones meet the needs of the applications they want to build. This includes both the information content and practical considerations of dependability (accuracy, availability etc), conditions of use, and commercial considerations. Sufficient information needs to be available within the ecosystem, via catalogues for example, to identify candidates and compare them.

Application developers then need to be able to arrange access to selected data resources

as simply as possible - for example, via uniform subscription processes and programming interfaces. Additional services such as development, test and hosting infrastructure may also be of interest.

Analytics service providers

Analytics service providers have started to emerge in the last few years, and naturally fit within an IoT ecosystem. They typically use specialist software tools to enable efficient analysis of particular kinds of data, for example very large datasets or high-frequency streams of data. Some providers offer general-purpose frameworks which an end user can exploit to define their own analytic workflows. Others, typically application domain experts, offer a more complete service, including consultancy, to end users who may not have the necessary data science expertise.

Platform providers

Platform providers have an enabling role in the IoT ecosystem. They do not directly

provide information or build applications but support stakeholders in other roles by providing a set of common services that all can use, allowing them to focus on their core concerns and accelerating innovation in the ecosystem. They may provide computing and storage infrastructure, offered for example as virtualised cloud services. More specific to IoT ecosystems, platform providers will offer the facility for information resources and value-added services to be advertised to potential users. Catalogues need to describe both the content and other attributes (quality measures, terms and conditions etc), ideally in a form that can be automatically processed. There is obvious benefit in a common approach to catalogues of information and services, which is a major motivation for the HyperCat initiative[1]. Platform services also include facilities for collecting, transforming, distributing and storing data from many independent and heterogeneous sources. The aim is to accept sensor data from any source, in any format and deliver it in a timely way to those wanting

to consume it, in a format convenient to the recipient. This approach logically centralises the management of data from many different sources and represents a natural integration point where various policies (for access control, messaging semantics, routing etc) can be managed and enforced.

As described above, information providers, application developers and analytics service providers may participate in the ecosystem in order to sell to other stakeholders. In that case the option to use commercially-focused services such as accounting, billing and payment, and customer support may be attractive.

End users

End users participate in the ecosystem by using information and applications made available to them. This might be an interactive application for people, viewed on a PC or mobile device, for example. On the other hand, it could be part of an automated system (e.g. building management) with no direct human involvement. In either case it is important that it delivers real value and a positive experience for the ultimate beneficiaries of the functionality provided by the other stakeholders. An IoT ecosystem will not be sustainable without the trust of its end users. For individual people, participation in the ecosystem is generally via an application. Often this involves information associated with their use of the application being collected, for example the user's location in the case of an app on a mobile phone.

The situation where the individual is an information provider needs to be addressed with care, particularly where personally identifiable or potentially sensitive information may be involved. Open engagement with end users which ensures they are properly informed and included in the ecosystem is essential.

MK:Smart

The MK:Smart project[2] is a collaborative initiative, partly funded by the Higher Education Council for England which is developing innovative solutions to support economic growth in Milton Keynes. An

important focus is the creation and support of an IoT ecosystem. In this example, the shared context is geographical – the city of Milton Keynes. The people who live and work there, organisations responsible for transport, energy and water supply, education institutions and businesses in Milton Keynes and the surrounding area, all have direct interest in similar information. They can all benefit from the availability of more up-to-date, reliable information from a variety of sources to understand their local environment and improve the way they live and work. The implicit connections being uncovered between previously separate application areas such as energy and water are of particular interest.

A focal point is the MK Data Hub, a system which aims to support the emerging local IoT ecosystem. It seeks to occupy the platform provider role described above, offering a number of services to other stakeholders in the project. It is characterised by a number of features which together represent a set of principles we regard as important for a sustainable IoT ecosystem focused on a community in a specific geographical area.

Data ownership is respected. People and organisations are assumed to retain ownership of their data, regardless of the fact that it is shared with others. Terms and conditions, including access controls, licences and pricing are defined by the data owner and enforced by the Data Hub. Access by other stakeholders is subject to accepting the terms defined for each data source.

Information is treated as a valuable asset. Where possible it is made open, but not necessarily free of charge or free of policy constraints required by the information provider. The ability to trade information is an important element so that an information provider can exploit its information assets commercially. Usage is metered at the individual data resource level and associated with user accounts. This gives the possibility of usage-based billing but also allows an analysis of usage patterns, giving an indication of what kind of information is found to be most widely used.

There is support for new sensor deployments. Sensor connectivity is made available with a range of cost effective solutions. Coordination and cooperation in the context of the MK:Smart ecosystem can identify opportunities to share network resources across multiple deployments. Application and service stores offer flexible business support systems so that developers can launch commercial applications quickly and easily.

An important feature is that the implementation of the MK Data Hub is designed to be secure and scalable and to address commercial concerns from the start, rather than just being a technical experiment. While there are a number of technical challenges, it is becoming clear that the ability to engage a wide range of stakeholders in an ecosystem and to understand how each participant benefits is the key to exploiting the potential of the Internet of Things.

Interoperability

The MK:Smart project aims, through real world case studies and applications, to illustrate the value of making information more freely available in improving current ways of working and enabling innovation. Demonstrating good practice, the benefits of making the best use of all available information and enabling new modes of cooperation between organisations are important contributions which are replicable in other situations.

We anticipate that there will be many IoT ecosystems, each focused on particular kinds of information. For example, they may be based on a geographical area, as in the MK:Smart project. The scope could be administrative (e.g associated with a city, region or nation) or environmental, such as a river, where flooding and drought can have significant impact on communities. Alternatively ecosystems may naturally develop around concepts such as health and wellbeing, or associated with a global supply chain, rather than a location. As well as the information content itself, individual IoT ecosystems may have specific requirements

on security, reliability and latency, particularly as IoT systems develop beyond sensors to include actuators which can directly affect the physical world. This diversity will clearly lead to distinct and varied technical solutions with a risk that opportunities for sharing potentially valuable information are lost. To avoid this, there is a need for consistency of approach and interoperability. While it is entirely reasonable for an individual or organisation to regard some or all of its data as private or confidential and to be unwilling to make it available to others, unnecessary technical barriers should be eliminated. Interoperability between ecosystems – based on interoperability between data hubs is therefore required.

HyperCat started as a way of achieving interoperability between multiple data hubs in different domains, based on a programme funded by InnovateUK[3] (the UK's innovation agency) in 2014, which involved eight separate IoT clusters. HyperCat is a specification for representing and exposing IoT catalogues using web technologies (including Uniform Resource Identifiers, Hypertext Transfer Protocol, JavaScript Object Notation etc) to make it easier to identify and discover data resources. If each data hub provides a consistent and uniform machine-readable catalogue, it is possible to create a view of resources across a set of hubs, allowing applications to find relevant data wherever it is held. The HyperCat specification is a pragmatic starting point, solving some important issues facing developers of IoT applications.

More generally, the ability for systems to be interconnected and to interoperate at greatly increased scale characterises the Internet of Things. While existing standards can already

be applied to IoT systems, there may be gaps and inconsistencies that need to be addressed. Many established standards development organisations, as well as industry consortia are now actively working on IoT.

AUTHORS' CONCLUSIONS

The Internet of Things is highlighting the importance of measuring and monitoring the physical world, the ability to know more about the state of individual systems and the environment within which they operate. There is a growing appreciation of the value of accurate timely information in effective decision making across many application areas. The importance of the relationships, implicit or explicit, between systems that have traditionally been managed separately is also becoming clearer. While there are many uncertainties and technical challenges to be faced as the IoT develops, the establishment of effective cooperative relationships between multiple stakeholders – organisations and individuals – is a significant contributor to the success of IoT applications.

A number of different roles within the ecosystem were identified, all of which need to derive clear benefit from participation. The MK:Smart project is exploring some of the practical issues associated with creating an IoT ecosystem, focused on the city of Milton Keynes. It is a broad collaboration engaging a variety of different stakeholders with the aim of solving some real world problems in a sustainable way, considering technical, commercial, organisational, and user experience issues. Initial experiences are positive and the project is expected to provide a tangible illustration of the potential of IoT ecosystems, and to identify good practice which can be replicated elsewhere. Interoperability is essential if the promise of IoT is to be fulfilled. There are already usable standards and specifications which can be applied to IoT systems. The increasing level of activity in standards development organisations and other bodies should identify and address the gaps specific to the new demands of the Internet of Things.

ABOUT THE AUTHORS

Dr John Davies
John is Chief Researcher in BT's Future Business Technology Research Practice. His interests include the application of semantic and IoT technologies. John leads BT's contribution to the HyperCat IoT interoperability project (hypercat.io) and was Project Director of the Stride smart transport collaborative project (stride-project.com), as well being involved in the MK:Smart smart city project (mksmart.org). He is currently chairman of the Web Intelligence, Mining and Semantics (WIMS-15) conference, is a Fellow of the British Computer Society and a chartered engineer and has published many books and papers in his areas of expertise.

Dr Mike Fisher
Mike is Chief Researcher in BT's Future Business Technology Research Practice. His interests include distributed systems and the convergence of computing and networks – most recently IoT and its applications. He has been engaged in a number of collaborative projects in these areas, including MK:Smart and Stride where he contributed to the HyperCat specification. Mike has had leading roles in a number of European collaborative projects, has been a member of European Expert Groups on Grid and Cloud computing and chaired the ETSI Technical Committee on Cloud. He is a member of the Institute of Physics and a chartered physicist.

FOOTNOTES

1 Hypercat. See: http://www.hypercat.io
2 MK:Smart. See: http://www.mksmart.org
3 InnovateUK. See: https://www.gov.uk/government/organisations/innovate-uk

ITP INSIGHT CALL
Want to talk to the authors?

To discuss this article and its content, join in the ITP Insight Call. Date to be confirmed at time of going to press.

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