Ubiquity Symposium

The Internet of Things

W3C Plans for Developing Standards for Open Markets of Services for the IoT

by Dave Raggett

Editor's Introduction

The Internet of Things (IoT) is being held back by divergent approaches that result in data silos, high costs, investment risks and reduced market opportunities. To realize the potential and unleash the network effect, [W3C](#) is focusing on the role of Web technologies for a platform of platforms as a basis for services spanning IoT platforms from microcontrollers to cloud-based server farms. Shared semantics are essential for discovery, interoperability, scaling and layering on top of existing protocols and platforms. For this purpose, metadata can be classified into: things, security, and communications, where things are considered to be virtual representations (software objects) for physical or abstract entities. Thing descriptions are modeled in terms of W3C's resource description framework (RDF). This includes the semantics for what kind of thing it is, and the data models for its events, properties and actions. The underlying protocols are free to use whichever communication patterns are appropriate to the context according to the constraints described by the given metadata. W3C is exploring the use of lightweight representations of metadata that are easy to author and process, even on resource constrained devices. The aim is to evolve the web from a web of pages to a "Web of Things."

Ubiquity Symposium

The Internet of Things

W3C Plans for Developing Standards for Open Markets of Services for the IoT*by Dave Raggett*

Continuing improvements in electronics are fueling the Internet of Things (IoT). An increasing range of devices is now available, although these often have very limited ability to interoperate with other devices and services. For the maker enthusiast, there is a range of controllers—such as Raspberry Pi and Arduino—and associated sensors and actuators. There are many possible application domains including consumer electronics (e.g. wearables and home automation), retail, manufacturing, construction, building management, transport, major utilities, healthcare, assistive technology, and scientific research.

Until recently most of the effort has been focused on the devices (sensors and actuators) and the communication technologies used to access them. Gartner expects IoT vendors to top \$263 billion in direct revenue by 2020, with most of that money deriving from services [1]. Web technologies are already being applied to the IoT, and it is now timely to consider how this could be expanded to help the IoT move beyond product silos into web-scale open ecosystems based upon open standards. This includes standards for identification, discovery and interoperation of services across platforms from different vendors, and will involve the need for rich descriptions and shared data models, as well as close attention given to security, privacy, scalability and accessibility. Open ecosystems will stimulate growth through the establishment of larger markets for developers and lifting the burden for tailoring products to vendor specific platforms.

There is plenty of potential for exploiting scripting languages like JavaScript, data encodings such as JSON and EXI, formats for data and metadata, including Linked Data, and protocols such as HTTP and WebSockets, to name just a few examples. JavaScript could be used for direct access to IoT sensors and actuators from the browser, in service platforms in the cloud or at the

network edge, and for device drivers in gateways that use IoT protocols to access embedded/constrained devices, and web protocols to expose them to service platforms.

Smartphones and tablets can be an attractive way to access IoT sensors and actuators, e.g. wearables and home automation. Smartphone operating systems like iOS and Android already offer APIs for native apps for access to embedded sensors and to nearby devices via Bluetooth or NFC. Work has started on defining standards for browser based APIs, e.g. the [W3C Web Bluetooth Community Group](#) and the [W3C NFC Working Group](#). Future work may cover additional protocols, e.g. CoAP and MQTT. Bluetooth Low Energy can be used to advertise devices to other devices in their neighborhood. Google is exploring this with their ideas for a “[Physical Web](#).” A similar approach is being developed by 3GPP for cellular networks (D2D proximity events), see Lin et al. for an overview [2].

Home hubs provide another opportunity for hosting services. Many people will have a home hub that functions as a Wi-Fi access point for broadband Internet access. It is reasonable to expect these devices to evolve into service platforms for accessing IoT devices around the home. For this to be really successful, the hubs will need to support open standards and allow users the freedom to install services from whomever they choose. There will also be a need for gateway devices that can be positioned around the home to reach IoT devices that aren't within range of the hub.

Cloud-based platforms can be designed to flexibly scale with the load on the platform. An example is the open source [Compose platform](#), which is based on top of Cloud Foundry. Service compositions are authored with [Node-RED](#) and then deployed to the Compose cloud. Service owners can define security policies, which are enforced through a mix of static flow analysis and dynamic monitoring. For an open market of services, there needs to be a standard way to describe the interfaces that services expose as well as the interfaces they in turn depend on. Linux package management provides a useful analogy, where packages have names and version numbers, and declare the packages names and version number ranges that they depend on to work. What is needed to encourage reuse of data and metadata vocabularies? One answer is to support repositories that developers can browse through to search for existing vocabularies. Likewise, developers need to be motivated to upload new vocabularies when the existing ones are not a good match. An example of such a repository is [schema.org](#).

What is needed to monetize services? This is obviously important for a healthy ecosystem of services. For web-scale ecosystems, we will need open standards that are independent of particular vendors. W3C has recently launched a [Web Payments Interest Group](#), which is seeking to decouple web applications from the means of payment. In principle, there could be upfront payments, subscription based payments, per use payments, or out of bound mechanisms, e.g. where particular users or groups of users are granted the rights to use a given service. Contract law is roughly similar across countries, and could form the basis for legally binding agreements between suppliers and consumers of services. This could pertain to payments and to data handling policies.

Cyber-physical systems are essentially control loops that bridge sensors and actuators to achieve system goals, e.g. traffic management on city streets, systems to maintain a comfortable environment in large buildings, and smart grids for electricity. The control can be expressed at multiple levels of abstraction, along with the means to distribute control to the network edge to address requirements for low latency and tightly coordinated synchronized control over multiple actuators. The protocols will need to be matched to fulfill the latency and jitter requirements, and there may be a need for the service layer to pass quality of service requirements down to the network layer. Latency may be subordinate to transactional robustness, and the requirements are likely to vary at different levels of abstraction. Cyber-physical systems create opportunities for pushing scripts to controllers along with the need for APIs for a wide range of protocols.

There is a huge variety of IoT protocols and these are continuing to evolve at a rapid pace. This suggests the need for abstraction layers that loosely couple to the underlying IoT technologies. This will simplify the development of services with less to learn and increase robustness in the face of changes to these technologies. Abstraction layers are also important for dealing with a heterogeneous mix of device vendors and versions.

IoT devices are often constrained and may not be software upgradeable. This puts them at risk when security flaws are identified. This can be worked around through the use of device gateways; the gateway itself will be software upgradeable, and can offer stronger security than is practical on hardware-constrained devices. In principle, gateways could use scripts for device

drivers. This raises the question of how to identify which driver is needed for a specific device, and the need for standard APIs for drivers to access the IoT protocols.

Identity is important for devices, users, applications and services, e.g. as part of end-to-end security and for trust management. Unlike regular web applications, we can't assume users are present and able to authenticate themselves. Trust management will entail the means to verify metadata, e.g. the provenance of data, the location of a given sensor, and so forth. This is analogous to knowing your customer requirements in the banking world. Trust can be built on known brands, strong vetting processes, and even crowd sourced reputation management.

Security will be an essential requirement for the “Web of Things.” With the increasing number of sensors all around us, we will need assurances of privacy from snooping and safety from attackers. Encryption of data is just one aspect, and there is a need for strong authentication of users, devices, services and applications. This will create opportunities for secure elements, hardware tokens and biometrics as supporting technologies.

With increasing dependence on systems based upon the Web of Things, we will need to pay careful attention to system resilience. Services will need to be able to cope with rapidly peaking demand loads. Back in 1994 when the web was very young, servers were unable to cope when large numbers of people wanted to see the new images of the collision of Comet Shoemaker-Levy with Jupiter, which produced fireballs the size of the Earth! We are now much more experienced with how to design server farms for scalability. Another challenge is dealing with a heterogeneous mix of device vendors and versions. As mentioned previously this can be addressed through abstraction layers and best practices. With so many devices, it is inevitable that some will fail, either through hardware faults or software bugs (e.g. botched upgrades). Services will need to be designed to be tolerant to failures, including the ability to spot when sensor readings are implausible and to employ appropriate workarounds. Cyber attacks by criminals and hostile states are a further challenge. This requires careful attention to fixing security flaws. The main principle is to employ defense in depth, along with monitoring and trip wires and the use of security zones.

Applications and services often need data at a higher level than the raw data provided by sensors. Moreover, data needs to be interpreted in the context of other sources of information. The same applies to control systems whose actions need to be translated in context into actions on lower level entities. The Web of Things needs to be able to model the real world at different levels of abstraction, and to enable open markets with free competition of services across these levels. The things in the Web of Things can be considered as virtual representations of objects.

A consequence of this is that the "things" in the Web of Things are not limited to connected devices, but can also include things that are not and cannot be connected such as people and places, and abstract ideas, such as events (e.g. a concert), organizations, and time periods (e.g. the 1970s). Each thing can have one or more virtual representations (avatars). Things can also have histories, e.g. for a car, recording the sequence of previous owners. Avatars have identities, rich descriptions, services, access control and data handling policies. Avatars have URIs (uniform resource identifiers) and are accessible via web technologies. Avatars make it easier to build applications and services that combine information from different sources and different levels of abstraction.

Open standards for services would further increase opportunities for smart, intent-based search. This starts in the usual way with typing in a search string. The search engine uses rules of thumb to recognize the intent and extract the associated search parameters, and pass them to registered services via the interface associated with this intent. These intent services can invoke other services as needed but have to be able to pass back the result in a fraction of a second for integration in the search engine's results page. Note the result could embed a more complex query that the user can then activate via tapping the link or icon. This allows for tasks which take substantial amounts of time to compute and which are delivered to the user as out of band notifications.

What are the legal implications for the Web of Things? One issue is around who accepts liability for errors or disruptions to services. For free or low cost services, this is likely to be the end user. For premium services, it may be the other way around, as determined by the contracts between service providers and consumers. For trust and delegation, it is likely there will be disputes around the vetting procedures. There are also likely to be upheavals for the insurance business as the Web of Things changes the likelihood of accidents.

To explore what is needed to realize the potential, the World Wide Web Consortium (W3C) held a workshop in [Berlin on 25-26 June 2014](#), hosted by Siemens. This has been followed up with the chartering of a [W3C Interest Group](#) for the Web of Things. The aim is to accelerate the development of open markets of applications and services based upon the role of web technologies for a combination of the Internet of Things with the web of data. The new interest group will start with a survey of use cases across industry sectors along with a study of existing solutions and standards relevant to the Web of Things. These surveys are essential for ensuring a shared understanding and will be a prerequisite for splitting work up into task forces that can then proceed in parallel. The most important success criterion for the Interest Group is whether it succeeds in establishing a consensus around a focus for the work in the W3C, defining what is needed that is in the W3C's scope, and driving that work into existing or new working groups as appropriate.

A separate topic for discussion is how advances in the Web of Things could enable humanity to overcome the challenges we currently face from resource depletion, climate change, and over population. Increasing automation at progressively higher level of skills would concentrate value production in the hands of a small minority. What does this mean for the rest of the population? How then do we create stable civil societies? This is getting away from computer science, but nonetheless is something we all need to think about in our increasingly unequal societies.

With Gartner predicting some 26 billion units for the IoT by 2020, it is likely that web is going to get a great deal bigger than it is now. We can look forward to a new phase of exponential growth like we saw in the early days of the web. In the process, new Internet giants will appear and old ones decline. We live in exciting times.

About the Author

Dr. Dave Raggett is a member of staff for W3C's European host (ERCIM) and a visiting professor with the University of the West of England. He has been closely involved with the development of core Web standards since 1992 contributing to HTML, HTTP and many others. He is currently leading W3C's efforts to realize the potential of the Web of Things.

References

[1] Gartner. (2014). Gartner Says 4.9 Billion Connected "Things" Will Be in Use in 2015. [Press release]. Retrieved from <http://www.gartner.com/newsroom/id/2905717>.

[2] Lin, X., Andrews, J.G., Ghosh, A., and Ratasuk, R. (2013). An Overview on 3GPP Device-to-Device Proximity Services. Retrieved from <http://arxiv.org/abs/1310.0116>.

DOI: 10.1145/2822531