

Building the Internet of Things

Jari Arkko, Ericsson Research

Abstract — *Everything that benefits from networking will eventually be connected. This is the basis for the interest in the “Internet of Things”, a reason why many research projects exist, why a large number of standards bodies want to build standards for it, and why there is a lot of commercial activity. However, on some areas, the Internet of Things pushes the limits of the current Internet. The Internet will evolve to meet the demand, as it has done in previous occasions. This paper argues that this evolution is more about enabling interoperability than about developing new technology. Most of the communications technology for creating the Internet of Things already exists, and the pain points are elsewhere.*

Keywords — Internet of Things, interoperability

I. INTRODUCTION

A key feature of the Internet is that different devices can work together: any browser works with any web server, almost all content is viewable by all devices, any device can plug into a home router, different networks can exchange routing information with each other, and so on. As the Internet has evolved, *interoperability* has always been a major concern, in terms of protocol design and extensibility, building products that in practice work well together with other devices, and setting standards.

In general, today's Internet builds on a key set of protocols that work extremely well between different types of devices and in varying types of networks: IP itself, TCP, DHCP, DNS, HTTP, TLS, HTML, XML, and so on. But in many cases there are still some components in the protocol stack that are proprietary, application-specific, available for limited platforms, or come from a single source. For instance, specific link layer technologies, content formats (Flash), or applications (Skype). This shows that it is important to balance the need for an interoperable Internet with the need to allow commercial innovation. Still, it is expected that over time, generally interesting components become available for all devices that need them.

This paper argues that we already have most of the technology that we need for building the Internet of Things, and that the problem is not so much about lack of technology but rather how to ensure that different pieces of equipment can work well together.

Section II discusses some of the existing technology that we believe is a key for building the Internet of Things. The rest of the paper discusses the concern of interoperability. Section III introduces the basic problem of a capability mismatch: how a small sensor may not be capable of using the same communication protocols as a full-fledged Internet host. Sections IV through VI highlight certain specific issues such as the

need for interoperability at a semantic level. Finally, Section VII concludes with some recommendations on how these concerns can be alleviated.

II. TECHNOLOGY FOR THINGS

There are many research proposals, ongoing projects, and standardization efforts in this area. It is perhaps important to emphasize that this does not mean we have to wait for the completion of all this work before we can start deploying our Internet of Things networks. Some of the ongoing work addresses important requirements in specific situations, some brings interoperability where we previously had none, some develops useful optimizations, some is interesting research, and some is merely exploring new designs.

But if we just look at what is already being deployed in real-life, it becomes clear that to a large extent, the Internet of Things technology is already here. In 2004, the utility company that provides energy for my house gave me a meter that uses cellular modem to upload information to a server in real-time. This was a standard solution for all new subscribers in Finland already back then. Going beyond the personal and national anecdote, the industry by and large is already deploying this technology. Utility companies with millions of metering devices, service companies with innovative ideas, health and sport related devices (weight scales that employ wireless LAN, for instance), e-book readers, tablets, cameras, and other gadgets with network connectivity, traffic applications that employ communications and positioning technology, building and surveillance solutions that run on top of IP, gateways that link legacy networks to the Internet, and so on.

Some of the key tools in the Internet of Things toolbox include:

- IP – including IPv6. It will also be necessary use the various mappings how IP runs over particular link layers (such as 6LOWPAN [1]). We also need the necessary tools that allow IPv6-based Things to communicate over legacy IPv4 networks (such as NAT64 [2]).

- Basic web technology: TCP, HTTP, HTML, XML. While as not optimized as some newer solutions, this is very easy to use, efficient when properly used, and guaranteed to pass through any home gateway firewall.
- Link layer technology such as cellular, wireless LAN, and ZigBee.
- User interfaces based on the web, and in some cases on SMS, e-mail, chat, or social media interfaces.

A more in-depth discussion of the available Internet protocol tools can be found in [3].

III. INTEROPERABILITY

Today's Internet is primarily characterized by applications with a human in the loop. A successful Internet application is one where the desired human experience is achieved. For instance, the desired visual effect is correctly rendered on screen. This makes interoperability a bit easier, as the humans are responsible for processing the “semantic” part of the communications. Today's Internet also consists of a relatively homogeneous set of devices. While there are differences between a smartphone, a laptop, and a high-end server, for instance, they are all still high power computing devices.

Some of the requirements and expected usage patterns in the Internet of Things will cause interoperability challenges. For instance, there is

- a capability mismatch between traditional Internet hosts and small devices,
- widely differing communication and processing bandwidths in different devices,
- needs for interoperability at a semantic level,
- different internetworking protocol choices (legacy vs. IP vs. IPv6), and
- solutions that are suitable for only some networks.

The two first items are a key problem. The desire to build large numbers of small, battery-operated, and inexpensive devices drives the need for simple solutions. Often these devices are not easily software upgradable, and their protocol and application suite is limited. Some of the typical limitations include:

- MTU limitations,
- simplified web protocols (COAP/UDP [4] instead of HTTP/TCP),
- single-stack instead of dual-stack,
- limited or no support for security that would be suitable for operation over the Internet,

- sleep schedule that does not allow for communication at all times,
- and so on.

These limitations would have no effect if the device only communicated to other similar devices, but they do have an effect when attempting to provide *Internet-wide interoperability* to such devices. For instance, clients that today employ HTTP would be unable to communicate with such a device. We believe that Internet-wide interoperability is required, as the system of connected devices usually consists of sensors, actuators, user interfaces, servers, and other components. Many of these components are expected to be devices in the traditional Internet. For instance, it is likely that computers and smartphones are used as the user interface for controlling many Internet of Things applications.

It is important to note that some of the capacity requirements would preclude direct communication to an Internet of Things device even if implemented exactly the same protocol stack as other devices in the Internet. For instance, a sensor whose value is interesting to a large audience may not be able to accommodate all requests.

IV. SEMANTIC INTEROPERABILITY

Most Internet applications designed for humans often require only transport of data from one place to another, and an accurate rendering of that data on the screen. It is not necessary to process or understand the data in any semantic manner.

Much of the current focus in the Internet of Things is also on the lower parts of the stack: designing the wireless networks, running IPv6 over them, getting routing to work, and using UDP/TCP and COAP/HTTP.

It is important to realize that this is **not** enough for true interoperability. For instance, it would not be enough for a light switch from one vendor to control lights from another. For true interoperability we need *semantic interoperability*, the ability of the devices to understand what the data they communicate **means**. Most often this would imply standardizing not just the protocols and data formats, but also the meaning of the data, e.g., that “1” in a particular field means that the light should be switched on. Standardizing the meanings is difficult and time consuming, however. It has to be done on a per-application basis and with application specific expertise.

There are of course different ways of achieving semantic interoperability. This does not always involve standards. Devices could accept program code that performs the required actions. For instance, a light

switch might accept a program fragment from a light bulb to run the user interface necessary to control the light. This is similar to how Flash-based applications can support new video codecs without requiring support from the browser or any Internet-wide agreement about the new coding format. It remains to be seen if programmable control models become popular in the Internet of Things.

Nevertheless, there should be some way for the light switch and the light bulb to agree how the lights are turned on. This is not to say that there is no benefit from an Internet of Things without it. There will always be a need for some proprietary or leading edge, non-standard communications. And even if none of the application layer communications would interoperate with each other, we would still have a common backbone for the Internet of things, consisting of the IP layer, routing, COAP/HTTP proxies, and so on. We call this the *Internet of Things transport network*. This would be tremendously valuable. But it would not enable an Internet of Things where any light works with any switch or any energy meter works with any provider's server.

V. AUTHORIZED INTEROPERABILITY

There are a number of security related challenges as well. Many of these fall into the capabilities category. But there is another, more fundamental issue. It is not enough that two endpoints support the same security mechanisms. The communicating parties also have to share some type of relationship that allows them to *authenticate* each other and *authorize* whatever actions are taking place. There are many ways to implement this, for instance with shared secrets, trusted third parties, or certificate infrastructures. It is relatively straightforward to set this up in small networks or within a single organization. Setting this up in a larger scale or in situations that require multiple participating organizations is going to be harder. For instance, home owners, manufacturers, and electricity utility companies might all want to control a particular home appliance.

VI. NETWORK-SPECIFIC SOLUTIONS

The Internet of Things is pushing the limits of technology in many areas. As we approach those limits we need to apply optimizations and design techniques to make our technical solutions feasible. But at the same time this may make our solutions less general than we would wish. For instance, the RPL routing protocol [5] has two modes optimized for different types of networks. Those modes are necessary, because without them it's not possible to support some important applications. However, the modes are incompatible and highly optimized implementations are unlikely to support both. As a result,

interoperability is not assured merely through the use of the same protocol. Note that while we use the two modes from RPL as an example, many similar issues exist elsewhere as well (different header compression types in 6LOWPAN [6], XML vs. JSON vs. binary XML for sensor data, and so on).

VII. CONCLUSIONS

Employing IP (and IPv6 in particular) for the Internet of Things is a necessary step. However, it is only the first step in ensuring a truly useful Internet of Things where different objects seamlessly communicate with each other. Some of the key areas where further work is needed include, for instance, standardization of application specific messages and semantics, and ensuring that each individual protocol specification is interoperable in all situations.

Looking back at the development of the Internet, one of the lessons that we can draw from it is to ensure that we have sufficiently general mechanisms that address most needs. Highly optimized and specialized solutions have rarely succeeded. Robustness and generality are often more important than mere performance. Based on this it is likely that most networks will actually employ pretty much standard Internet technology as we already know it today: IP, web protocols, and existing link layers. We need to be careful about spending too much effort in optimizations with narrow usability. Achieving interoperability has been a far more important success criteria for the Internet, and this is where we should spend our future efforts.

REFERENCES

- [1] G. Montenegro, N. Kushalnagar, J. Hui, D. Culler. Transmission of IPv6 Packets over IEEE 802.15.4 Networks. RFC 4944, IETF, September, 2007.
- [2] F. Baker, X. Li, C. Bao, K. Yin. Framework for IPv4/IPv6 Translation. Internet Draft draft-ietf-behave-v6v4-framework (work in progress), IETF, August, 2010.
- [3] F. Baker, D. Meyer. Internet Protocols for the Smart Grid. Internet Draft draft-baker-ietf-core (work in progress), IETF, March, 2011.
- [4] Z. Shelby, K. Hartke, C. Bormann, B. Frank. Constrained Application Protocol (CoAP). Internet Draft draft-ietf-core-coap (work in progress), IETF, March 2011.
- [5] T. Winter, P. Thubert, A. Brandt, T. Clausen, J. Hui, R. Kelsey, P. Levis, K. Pister, R. Struik, JP. Vasseur. RPL: IPv6 Routing Protocol for Low power and Lossy Networks. Internet Draft draft-ietf-roll-rpl (work in progress), IETF, March 2011.
- [6] J. Hui, P. Thubert. Compression Format for IPv6 Datagrams in Low Power and Lossy Networks (6LoWPAN). Internet Draft draft-ietf-6lowpan-hc (work in progress), IETF, February, 2011.



Jari Arkko is a researcher with Ericsson Research in Jorvas, Finland. He serves currently also as one of the Area Directors at the Internet Engineering Task Force (IETF). His main interests include Internet Architecture, IPv6, the Internet of Things, and social media. He frequently communicates with his toaster on Facebook.