

The Things in the Internet of Things

Stephan Haller

SAP Research Center Zurich

stephan.haller@sap.com

Abstract— The Internet of Things is a hyped term and many definitions for it exist. Worse still, it comes with a lot of related terminology that is not used uniformly either, hindering scientific discourse. This paper tries to bring clarity by describing the most important terms like things, devices, entities of interest, resources, addressing, identity and, more importantly, the relationships between them.

Keywords- Internet of Things; Devices; Resources; Services; Addressing; Identification; Resolution; Discovery

I. INTRODUCTION

There currently different opinions on what the Internet of Things is, ranging from the original concept proposed by Kevin Ashton at the Auto-ID Center [1] of an informational network that allows the look-up of information about real-world objects by means of a unique ID called Electronic Product Code (EPC) and a resolution mechanism (ONS), to a network of sensors, actuators and autonomous objects interacting with each other directly [2]. Machine-to-machine (M2M) communication is another term sometimes associated with the later [3]. With so many diverging opinions, and with the hype that has evolved around the term, it is not surprising that there are equally many attempts at providing a definition of the Internet of Things [4][5][6].

More important though than agreeing on a definition of the overall term is to have a common understanding of the components and concepts that constitute the Internet of Things. First and foremost, what are the things in the Internet of Things? What is the relation to devices, resources and services? Second, what do we mean when we talk about addressing, identification and resolution? Looking at discussions at conferences and within international research projects, one can see that these terms are used with different meanings by different people, and that the terminology is often mixed up, leading to confusion and hindering scientific discourse. The goal of this paper is to bring some clarity into these discussions. It can also be seen as a call to a common usage of the terminology associated with the Internet of Things.

II. THINGS, DEVICES AND RESOURCES

All the different definitions of the term “Internet of Things” have in common that it is related to the integration of the physical world with the virtual world of the Internet. There are physical objects one wants to be able to track, to monitor and to interact with. Examples include inanimate objects like pallets, boxes containing consumer goods, cars, machines, fridges –

and maybe even the infamous carton of milk or cup of yoghurt – as well as animate objects like animals and humans. These are the things of the Internet of Things – or to use a clearer term, the *entities of interest* [7]. Buildings, rooms and things in the environment like rivers and glaciers can also be entities of interest. Basically any object including the *attributes* that describe it and its state that is relevant *from a user or application perspective* can be regarded as an entity of interest.

In order to monitor and interact with one or more entities and make the connection to the Internet, technical communication *devices* are required. The devices can be attached to or embedded in the entities themselves – thus creating smart things –, or they can be installed in the environment of the things to be monitored. Typical examples of devices include RFID readers, sensors and actuators, embedded computers as well as mobile phones. While there is a school of thought that regards the devices as the things in the Internet of Things, such an approach seems too limited, as businesses and consumers are more interested the physical objects rather than any technical devices needed for monitoring and communication. Having said that, it needs to be noted that devices constitute entities of interest in its own right when looking at them *from a technical or management perspective*. Thus, devices are a subset of all the things in the Internet of Things. However, for reasons of clarity this case where the thing, the device and the entity of interest are the same should be treated as a special case.

Devices usually host *resources*: These are computational elements that provide the technical link to the entities of interest – e.g., they offer information about the thing, like an identifier or sensed data, and they may provide actuation capabilities as well. Access to resources from the outside world finally happens through *services*. Resources may offer a service interface directly, or services inside the network act as proxies for the actual resources, possibly providing additional levels of aggregation and abstraction. RESTful services [8][9] can be used and are most appropriate when accessing resources directly, but other implementation technologies like SOAP [10] or Device Profile for Web Services (DPWS) [11] are also possible; in particular higher-level aggregated services that have to be integrated with enterprise applications. When using REST, the distinction between resource and service becomes blurry, but it can be disambiguated in the following way: use the term service when focusing on the application integration and accessing aspects, and talk about resources when looking from more low level component and deployment perspective.

The relationship between all these terms is schematically summarized in Fig. 1: An entity of interest is monitored by a device in the environment, or it can also have a device attached




Figure 1: Relationship between things, devices, resources and services

to or embedded in it. As described above both classes of devices can be seen as entities of interest when looking from a management perspective, hence the subclass relationship. The device hosts one or more resources which are accessed through services.

III. EXAMPLES OF THINGS AND DEVICES

The distinction between entities of interest and devices is often clear – the entity of interest is the object that has some value for the observer, the device is a technical component needed to observe or interact with the entity of interest. There are however many cases where the distinction is more difficult and therefore justifies to be elaborated upon. In the following, a few examples in common application areas for the Internet of Things are discussed: Logistics, energy monitoring and management, as well as public safety and disaster management.

Let's start by looking at container filled with some temperature-sensitive chemical that is transported from A to B. Its location needs to be tracked continuously, and for quality assurance reasons, the temperature must be monitored to ensure that it is within certain boundaries at all times. To do this, a wireless sensor tag is attached to the container. The tag provides both communication as well as temperature recording capabilities. In this simple example the distinction is obvious: the container is the entity of interest, and the sensor tag is the attached device¹.

But what about a pallet with a simple barcode label attached? The pallet is clearly an entity of interest from a supply chain application perspective, but the barcode label? One could argue that as the label is an artificial component attached to the pallet in order to be able to identify and track it, it also qualifies to be called a device. However, that would be stretching the common understanding of the term device. It makes more sense to regard the label as a „feature“ of the entity of interest itself; and the tracking of the pallet is done with the help of barcode scanners which thus constitute environmental devices. To strengthen the point, if we would

look at the label as a device, then the imprinted barcode would qualify as a resource, which again would be stretching the common understanding of these terms. It looks different though when we replace the barcode label with a passive RFID tag: The tag with its electronic circuitry and the communication capabilities can reasonably be called a device, and its memory containing an ID and possibly additional user data also qualifies as a resource. RFID interrogators are – similar to barcode scanners – definitely environmental devices monitoring things that pass through their reading fields. Only from the perspective of the administrator of the interrogator network they are also entities of interest.

In the field of energy monitoring, smart meters to remotely monitor the consumption of electricity – and in smart grid scenarios, also distributed electricity production – are one of the key components to minimize energy usage and allowing an accurate billing based on individual consumption. In more and more countries the usage of smart meters is mandated [12]. They also constitute an interesting example for this paper, as they clearly are both a device and an entity of interest. When looking at their purpose, they have to be considered a device that monitors the energy consumption of, e.g., a household – the entity of interest. But they are complex enough that they constitute an entity of interest to the organization responsible for managing and maintaining them.

In more advanced energy management scenarios, smart meters are also used to shape the energy demand to avoid peak loads and thus reduce the overall carbon footprint. For example, a freezer in a household may be shut down for a certain time without any detrimental effects to its contents. Such a freezer is an example where the distinction between entity of interest and device is simple: The freezer is the entity of interest, and it contains a controller component (the device) that allows the freezer to be turned on or off remotely.

The last example considers micro unmanned aerial vehicles (MUAV) that are becoming increasingly popular in disaster management. For example, they can be used to monitor the spread and the contamination of specific areas after the release of some hazardous substances [13]. As such, they are monitoring environmental devices that observe, e.g., a building, a plant, or a city district. For the hazard response team, these latter are the entities of interest. But for the operator of the MUAV – flight control – the MUAV itself is very much the entity of interest.

IV. ADDRESSING, IDENTIFICATION AND RESOLUTION

Unfortunately one can often see a confusion between the terms identity and identifiers. Identity is a philosophical concept meaning „whatever makes an entity definable and recognizable“ [14]. An entity of interest only has one identity, but it might have several unique identifiers (ID's) associated with it. These ID's are used to disambiguate two things from each other, and depending on the context, different ID's may be used. An ID can be compared to a social security number or a key value for looking up records about the thing in a data base. Many types of different ID schemes have been proposed for IDs in the Internet of Things [15], and it is unlikely that we will have one common scheme across the globe and across industries.

¹ Note that the sensor tag is however also an entity of interest for someone who is managing all tags, e.g., monitoring their battery status.

An address on the other hand is a technical term for accessing – “talking to” – either a device or a service. In the case of devices, the ID and the address often are the same, e.g., an IPv6 or MAC address, but in general they are not. As an example, let’s have a look again at the container filled with a chemical. The container has an ID, e.g., a Serial Shipping Container Code (SSCC). This ID can be used to find in a data base information like the type of chemical currently in the container or its current location. The sensor tag attached however might have an IPv6 address that can be used to query the sensor for the current temperature. The temperature readings could then of course be saved in the data base as properties of the entity of interest and thus become accessible also via the SSCC again, but the point is that the ID links to properties of the entity of interest, while the address is directly used for communication with the device.

In order to find information about a thing with a specific ID, two approaches are possible: *Resolution* and *discovery*. Resolving a given ID leads to a set of addresses of information and interaction services. Information services allow querying, changing and adding information about the thing in question, while interaction services enable direct interaction with the thing by accessing the resources of the associated devices. Resolution is a straight-forward process based on a-priori knowledge that will yield at least one address that should have information about the object. For example, the standard resolution mechanism for EPCs, the Object Naming Service ONS, will return with addresses of EPC Information Systems (EPCIS) of the original manufacturer of the product. In a second step, these EPCIS can be queried to return information about the thing. What information – if any – is returned depends then on confidentiality policies and potentially successful authentication of the requester.

Discovery on the other hand is more like “googling” for information without a-priori knowledge, finding in the process previously unknown sources of information [16]. Authentication is usually part of the discovery processes, i.e., only addresses of information services that are willing to provide information are reported back.

V. CONCLUSION

In this paper, the most important terms in regards to the Internet of Things and the relationship between these terms have been explained: the things, devices, resources and services, as well as identification, addressing, resolution and discovery. It has been shown – in particular regarding the distinction between the entity of interest and the device – that an absolute, clear-cut categorization is not always possible. Rather, it depends on the perspective from which one looks at a particular thing.

The intent was to remove some of the hindrances in scientific discourse and the development of the Internet of Things as one key aspect of an overall Future Internet. Acceptance of these definitions and uniform use in the future would ensure that research and development on the topic of the Internet of Things can progress more easily.

Having clear terminology will allow to focus on the real research issues like how to connect and interact with a myriad of heterogeneous devices, how to deploy and manage such

infrastructures, and how to model business processes that interact with things in the real world. Only by solving these real issues will it become possible to actually reap the many potential benefits of the Internet of Things that have been proposed; be it in supply chain management, the energy grid, health care, environmental management or public safety.

ACKNOWLEDGMENT

I would like to thank Elgar Fleisch and Thomas Michael Bohnert for reviewing an early draft of this paper, their valuable input and the many fruitful discussions we had on this and related topics.

REFERENCES

- [1] S. Sarma, D.L. Brock, and K. Ashton, “The Networked Physical World, Proposals for Engineering the Next Generation of Computing, Commerce & Automatic-Identification”, Auto-ID Center White Paper, October 2000
- [2] M. Presser, P. Daras, N. Baker, S. Karnouskos, A. Gluhak et al., “Real-World Internet – Position Paper”, position paper of the Real-World Internet (RWI) Cluster of the Future Internet Assembly (FIA). http://rwi.future-internet.eu/index.php/Position_Paper, last accessed May 27, 2010
- [3] Beecham Research, “Business Opportunities from Remote Device Management: M2M and the Internet of Things”, 2008.
- [4] S. Haller, S. Karnouskos, and C. Schroth, “The Internet of Things in an Enterprise Context”, in J. Domingue, D. Fensel und P. Traverso (Eds.), “First Future Internet Symposium - FIS 2008”, LNCS 5468, Springer Verlag 2009, pp. 14-28.
- [5] O. Vermesan, M. Harrison, H. Vogt, K. Kalaboukas, M. Tomasella et al. (Eds.), “The Internet of Things - Strategic Research Roadmap”, Cluster of European Research Projects on the Internet of Things, CERP-IoT, 2009.
- [6] EU FP7 Project CASAGRAS, “CASAGRAS Final Report: RFID and the Inclusive Model for the Internet of Things”, 2009, pp. 10-12.
- [7] F. Carrez (Ed.), M. Bauer, T. Baugé, J. Bernat, M. Bohli et al., “SENSEI Reference Architecture”, EU FP7 Project SENSEI, Deliverable D3.2, 2009.
- [8] L. Richardson and S. Ruby, “RESTful Web Services”. O’Reilly Media, May 2007.
- [9] D. Guinard, V. Trifa, T. Pham, and O. Liechti, “Towards Physical Mashups in the Web of Things”, in Proc. of INSS 2009 (IEEE Sixth International Conference on Networked Sensing Systems), Pittsburgh, USA, June 2009
- [10] World-Wide Web Consortium (W3C), “SOAP Version 1.2”, W3C Recommendation, 2nd edition, 27 April 2008
- [11] E. Zeeb, A. Bobek, H. Bohn, S. Prüter, A. Pohl, H. Krumm, I. Lück, F. Golasowski, and D. Timmermann, “WS4D: SOA-Toolkits making embedded systems ready for Web Services,” in Proceedings of the Open Source Software and Product Lines Workshop (OSSPL07), 2007.
- [12] Berg Insight, “Smart Metering in Western Europe”, M2M Research Series, June 2009
- [13] K. Daniel, B. Dusza, A. Lewandowski, and C. Wietfeld, “AirShield: A System-of-Systems MUAV Remote Sensing Architecture for Disaster Response”, IEEE International Systems Conference, March 2009, pp. 196-200.
- [14] Wikipedia, “Definition of Identity (philosophy)”, [http://en.wikipedia.org/wiki/Identity_\(philosophy\)](http://en.wikipedia.org/wiki/Identity_(philosophy)), last accessed May 22, 2010
- [15] EU FP7 Project CASAGRAS, “CASAGRAS Final Report: RFID and the Inclusive Model for the Internet of Things”, 2009, pp. 43-54.
- [16] C. Kürschner, C. Condea, and O. Kasten, “Discovery Service Design in the EPCglobal Network : Towards Full Supply Chain Visibility”, in C. Floerkemeier et al. (Eds.), “Internet of Things – IOT 2008”, LNCS 4952, Springer Verlag 2008, pp. 19-34.