

Information Models for an Interoperable Web of Things

Michael Koster*

*ARM Limited, Email: Michael.Koster@arm.com

I. ABSTRACT

A Web Of Things will enable application level interoperability. Application software will make use of sensors, actuators, and other data sources, analogous to the way humans make use of information on the World Wide Web. To make this a reality, hypermedia driven APIs can be annotated with machine understandable hyperlinks. Standards developed in the IETF CoRE (Constrained RESTful Environments) working group provide for standard web object encapsulation of sensor data with appropriate semantic linking and protocols for web scale resource discovery. We are seeing the emergence of architectures for the deployment of IoT services. What is still needed for application interoperability is a common set of high level patterns for the creation, manipulation, and use of descriptive, machine understandable hyperlinks, built on underlying web standard mechanisms and protocols. Information models¹ can be built on a base of reusable concepts, relations, and attributes.

The W3C could become a venue for standardizing information models that accommodate existing information and data models, enhance the functionality of existing models, as well as standardize new information models.

II. A WEB OF THINGS

The World Wide Web enables anyone to access to any information on any server using any web browser, independent of operating systems, enabled by a core set of protocols consisting of the IP, TCP, HTTP and HTML.

The Internet Of Things results from being able to connect everyday objects to software using networks and internet protocols [2]. This provides an infrastructure for scalable interconnection using the existing internet and new sensor networks.

The next evolution will be to enable application software to understand data from sensors, actuators, and other data sources in the way that people using web browsers understand information on the World Wide Web. The coming Web Of

Things will be enabled by application-level interoperability, that is interoperability between application software and network connected things. As with the WWW of information, the Web Of Things will connect any application to any connected thing using a choice of underlying transport protocols.

III. THE MACHINE INTEROPERABILITY PROBLEM

Humans are able to use web browsers to access information using the relatively simple mechanisms of hyperlinks associated with pictures, text, and other visual cues embedded in web pages and infographics. This is enabled by our ability to rely on visual metaphor, text association, spatial cues, and other high level human cognitive function.

Software needs to be able to access information from sensors using similar simple mechanisms, but software cannot rely on cognitive ability and visual metaphor to select the proper links. What is needed is a higher level of information which includes descriptive attributes sufficient to enable software to automatically discover and link to the appropriate resources.

The existing mime-types are sufficient for web scale interaction, but another layer of metadata is needed to create a machine understandable resource discovery and linking language.

IV. WWW ARCHITECTURE FOR THINGS AND SERVICES

The World Wide Web has evolved from its origins in simple HTML, to a rich set of protocols and abstractions that are responsible for the operation of the web at scale. A more recent development in web architecture is the adoption of the REST architecture style. There is a huge advantage provided by REST in the externalization of application state. This enables stateless applications and is key to achieving web scale.

The disadvantage of REST in dynamic systems is the lack of an architecturally well defined asynchronous, or "push" notification and update mechanism. In practice this is not a problem as there are many mechanisms available, including multi-part GET responses [3], using PUT or POST as commonly done in REST APIs to update endpoints, binding to a message protocol like MQTT [4] or XMPP [5], and use of the bidirectional socket connection [6].

REST endpoints are referenced by URLs, enabling hyperlinks to be used by application software to discover and interact with resources. Such hypermedia driven APIs enable the decoupling of connected things from dedicated services and applications, enabling broad interoperability.

POSITION PAPER FOR THE 'W3C WORKSHOP ON THE WEB OF THINGS - ENABLERS AND SERVICES FOR AN OPEN WEB OF DEVICES', 25 - 26 JUNE 2014, BERLIN, GERMANY.

¹There is a lot of terminology confusion between information models and data models. For this purpose RFC 3444 [1] has been published to provide guidance. As a summary, the main difference between the two terms can be described as follows. The main purpose of an Information Model is to model managed objects at a conceptual level, independent of any specific implementations or protocols used to transport the data. Data models are defined at a lower level of abstraction and include many details intended for implementors and include protocol-specific constructs.

V. INFORMATION MODELS

Application software can automatically discover and link to resources using link metadata, that is metadata embedded within hyperlinks. These hyperlinks have the form of a subject URL pointing to a resource and a collection of relation/attribute pairs. Software examines the relations and attributes in order to select links to suitable resources. Examples of metadata links are found in RFC 6690 [7] and Hypercat [8].

Also known as semantic hyperlinks, there can be many links and relations for a given resource. This allows rich descriptions to be created, with layers of context and resource bindings added as needed. The resulting scheme only requires a subset of general web linking, and a subset of the Semantic Web. The focus is on abstraction and virtualization of physical things for the purpose of monitoring and control.

Information models can be built from collections of semantic hyperlinks. These models describe the information provided by the linked resource sufficiently well to enable application software to correctly select the appropriate resources.

VI. WEB OBJECTS, RESOURCE DIRECTORIES, AND CATALOGS

Web objects encapsulate observable data from sensors, semantic links, and other metadata at an endpoint pointed to by a URI. A web object may map to a physical object, or a composite web object constructed of resources from other objects or data sources.

A web object may encapsulate a set of resources and their associated metadata, where each resource is a distinct URL from the base URI. There may be a common structure to the data and metadata endpoints, for example the well-known/core resource in CoAP objects that contains semantic links.

An example of an IPSO [9] web object, an information model for a temperature sensor, is shown in Table I and in Table II. IPSO objects are structured according to the Open Mobile Alliance (OMA) Lightweight Machine-to-Machine (LightweightM2M) object template [10].

TABLE I
EXAMPLE: IPSO OBJECT (TEMPERATURE)

Object	Object ID	Object URN
IPSO Temperature	3303	urn:oma:lwm2m:ext:3303

TABLE II
EXAMPLE: IPSO RESOURCE INFO (TEMPERATURE)

Name	Resource ID	Access Type	Data Type	Units
Sensor Value	5700	read	Float	Celsius
Min Measured Value	5601	read	Float	Celsius
Max Measured Value	5602	read	Float	Celsius
Min Range Value	5603	read	Float	Celsius
Max Range Value	5604	read	Float	Celsius

In this example, the Sensor Value resource is the measurement data, and the other resources are metadata. Here, the Sensor Value would be referenced by constructing the example URL: `coap://example.com:5683:/sensors/3303/0/5700` where 3303/0 refers to temperature object type, instance 0, and 5700 points to the Sensor Value resource within the object.

In this example, the object metadata are exposed as resource endpoints in addition to the Sensor Value measurement data. Semantic hyperlinks may also be encapsulated with a web object, as described in the ./well-known/core interface in RFC6690 [7]

To make the objects discoverable by applications, the object's semantic links need to be accessible by the application. One way is to allow the application to directly read the object's encapsulated metadata. More likely, there will be a middleware platform that the object's semantic links can be uploaded to, which is also reachable by the application software. An example is the CoRE Resource Directory [11], which provides capabilities for registering device metadata to make it available to applications for discovery by relation and attribute. Resource catalogs are another name for similar functionality, see [8].

For example, a sensor device may be shipped with metadata stored internally. When the device is installed, it can discover a nearby Resource Directory server and register itself and its metadata with the server. Application software and software on other devices can then discover the new sensor device and its attributes by sending queries to the Resource Directory.

VII. HYPERMEDIA DRIVEN API COMPOSITION

The information model for an object can contain information about the objects structure and resource types. This can enable APIs to automatically be composed by middleware and automatically consumed by application software.

Additional metadata can indicate context, such as geographical location, and bindings, such as message protocols and event handlers, as well as access control information.

Applications can use templates based on models to facilitate discovery and linking to resources by attribute and relation, including context and bindings. Consider the following examples:

- 1) Calculate average temperature within a region by linking to all of the resources that expose temperature measurements and are within a geographically bounded region.
- 2) Turn on a light in the room I am presently occupying.

Applying a filter based on a template would return objects and resources that satisfy the relations, as well as additional metadata that specify properties, for example engineering units to enable auto-conversion. An example can be found in Section 7 of [11].

VIII. THE ROLE OF THE W3C

Existing standards provide an infrastructure for connecting Things to the Internet. What is needed are information models for diverse application domains. Although standardization efforts on defining information and data models have been started in other fora already, the W3C might offer a unique

venue for defining additional information models by involving a wider community. The author hopes that the workshop will offer a possibility to discuss the need for information model standards for Web of Things, scope and timeframe.

REFERENCES

- [1] A. Pras and J. Schoenwaelder, "On the Difference between Information Models and Data Models," Internet Engineering Task Force, RFC 3444, Jan. 2003. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc3444.txt>
- [2] H. Tschofenig and J. Arkko, "Report from the Smart Object Workshop," Internet Engineering Task Force, RFC 6574, Apr. 2012. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc6574.txt>
- [3] Z. Shelby, K. Hartke, and C. Bormann, "Constrained Application Protocol (CoAP)," Internet Engineering Task Force, Internet-Draft draft-ietf-core-coap-18, 06 2028, work in progress. [Online]. Available: <http://www.ietf.org/internet-drafts/draft-ietf-core-coap-18.txt>
- [4] IBM, "MQ Telemetry Transport (MQTT) V3.1 Protocol Specification," Aug. 2010. [Online]. Available: <https://www.ibm.com/developerworks/library/ws-mqtt/>
- [5] P. Saint-Andre, "Extensible Messaging and Presence Protocol (XMPP): Core," Internet Engineering Task Force, RFC 6120, Mar. 2011. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc6120.txt>
- [6] I. Fette and A. Melnikov, "The WebSocket Protocol," Internet Engineering Task Force, RFC 6455, Dec. 2011. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc6455.txt>
- [7] Z. Shelby, "Constrained RESTful Environments (CoRE) Link Format," Internet Engineering Task Force, RFC 6690, Aug. 2012. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc6690.txt>
- [8] "HyperCat," Sep. 2013. [Online]. Available: <http://wiki.1248.io/doku.php?id=hypercat>
- [9] "IP for Smart Objects (IPSO) Alliance," May 2014. [Online]. Available: <http://www.ipso-alliance.org/>
- [10] OMA, "OMA LightweightM2M v1.0," Dec. 2013. [Online]. Available: http://technical.openmobilealliance.org/Technical/release_program/lightweightM2M_v1_0.aspx
- [11] Z. Shelby, C. Bormann, and S. Krco, "CoRE Resource Directory," Internet Engineering Task Force, Internet-Draft draft-ietf-core-resource-directory-01, 12 2011, work in progress. [Online]. Available: <http://www.ietf.org/internet-drafts/draft-ietf-core-resource-directory-01.txt>