



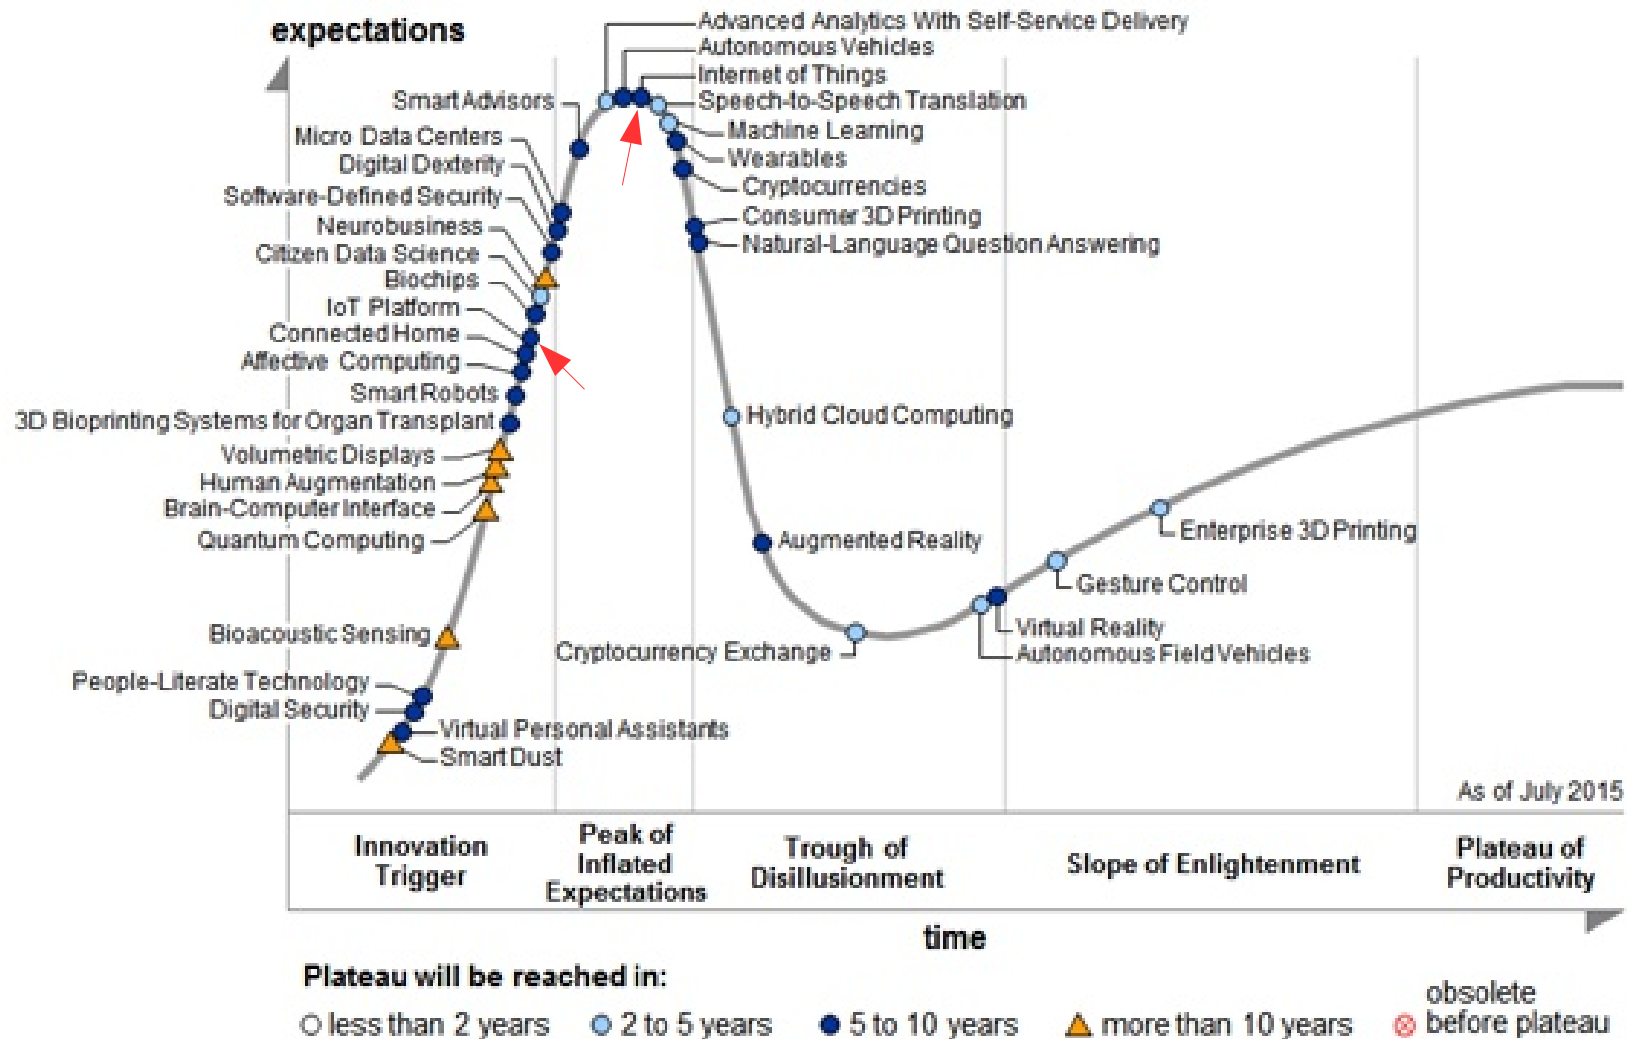
Using semantics and rich metadata to bridge IoT silos — W3C's work on the Web of Things

Dave Raggett <dsr@w3.org>

ETSI M2M Workshop

9 December 2015

Now is the time to work on IoT standards



* Gartner's Hype cycle: The IoT remains at the peak for yet another year and is still very immature



- W3C's mission to lead the Web to its full potential
- Founded by Sir Tim Berners-Lee inventor of the Web
- W3C is a member funded international organisation focusing on developing standards for Web & semantic technologies
 - e.g. HTML, CSS, scripting APIs, XML, SVG, the Semantic Web and Linked Data, ...



Challenges

- **Discovery** of services
 - The benefits of a *lingua franca*, and its limitations
- **Composition** of services
 - From different vendors for an open market of services
- **Monetization** of services
 - Support for a wide variety of models
- **Security**, privacy, safety, compliance, resilience
- Scaling on multiple dimensions
 - Scaling across **devices** from microcontrollers to massive cloud based server farms
 - Scaling across **platform and services** from different vendors and built upon different standards
 - Scaling across application **domains**
- The inevitability of evolutionary change in complex ecosystems
 - Weakly coupled communities will evolve independently
 - How to support “trade” across these communities

Helping not hindering convergence

- The IoT suffers from **fragmentation**
- Adding yet another platform will make this worse
- W3C wants to avoid this and instead make it easier for developers to create services that span platforms and enable an open market of services
- Hence a focus on a **platform of platforms**
- An abstraction layer that embraces existing platforms
- This will rely upon **rich metadata standards**
- There are many Industry Alliances and SDO's
 - We all need to **work together** to realise the potential
 - Open standards that reduce the cost of development, reduce the risk to investors, increase the market opportunities

Web of Things


- Making life easier for application developers by providing a **simple scripting** model
 - Things standing for physical and abstract entities
 - Things as software objects with **properties, actions** and **events**
 - Applications scripts **decoupled** from underlying protocols which can be selected according to need
 - **Servers can further choose** which **communication patterns** to use, e.g. push, pull, pub-sub and peer to peer as appropriate
 - Potential for **multiplexing data** from multiple sensors
 - Potential for **buffering sensor data** (optimise battery, network)
 - Dealing with **battery operated devices that sleep a lot**
 - Based upon **shared semantics and rich metadata**
- Server creates a software object based upon the thing's description
 - What properties, actions and events does it have?

From the Web of Pages to the Web of Things

- The **fundamental** elements of Web architecture
 - **Addresses, Resources, Protocols**
- For the Web of Pages
 - URLs, HTML and HTTP
- For the Web of Things
 - URIs, Thing descriptions, a suite of protocols
 - Different protocols are suited to different contexts
- Declarative formats for resources allow search engines to spider the Web and create indexes
 - HTML – Hypertext Links
 - **Thing descriptions** – Links between dependent “things”

Web Servers at Many Scales

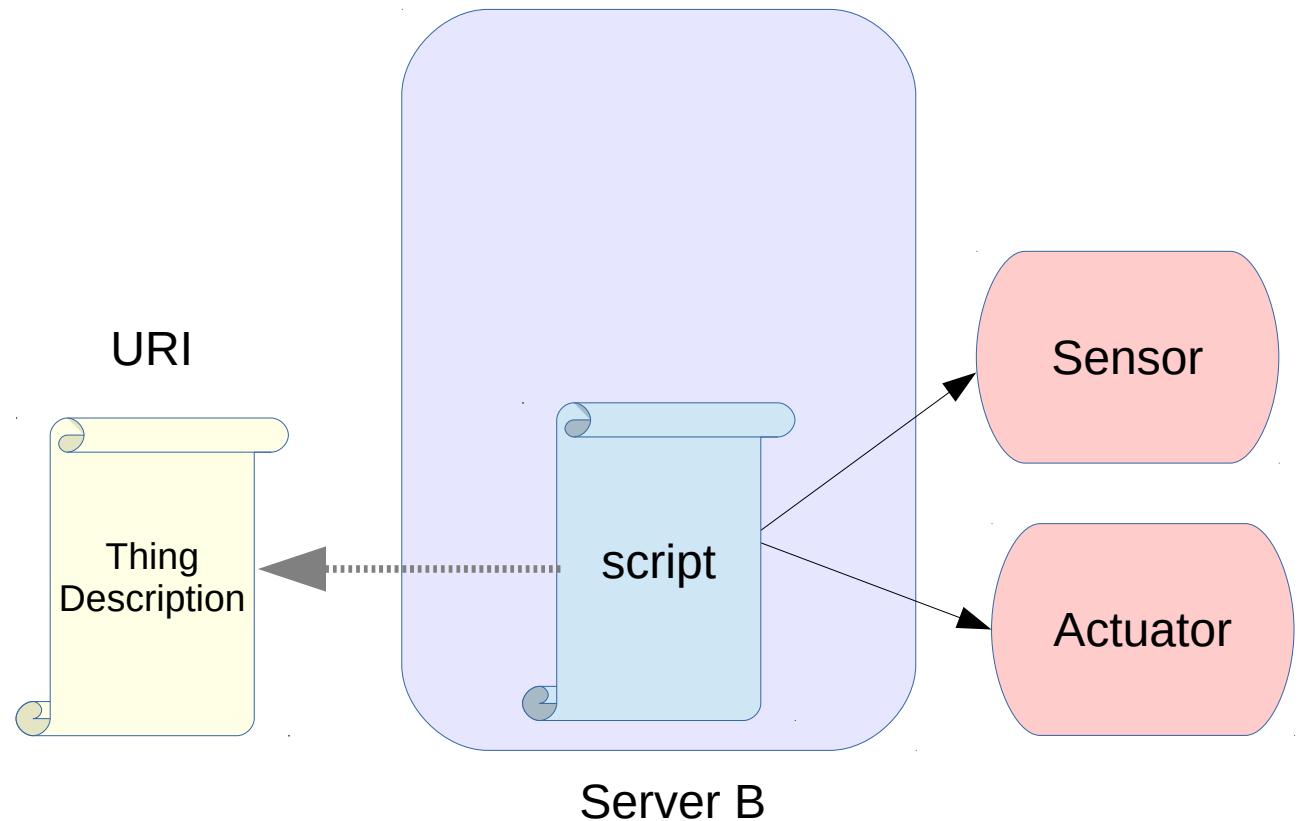
Web of Things servers can be realised at many scales from microcontrollers to clouds



Servers are free to choose which scripting languages they support
Precompile service behaviour for constrained devices


Distributed Web of Things

- Virtual representations of physical or abstract entities for use by application scripts
 - Each thing has a URI for a description which is used by the server to create a software object for applications scripts to interact with in terms of the thing's properties, actions, events and metadata



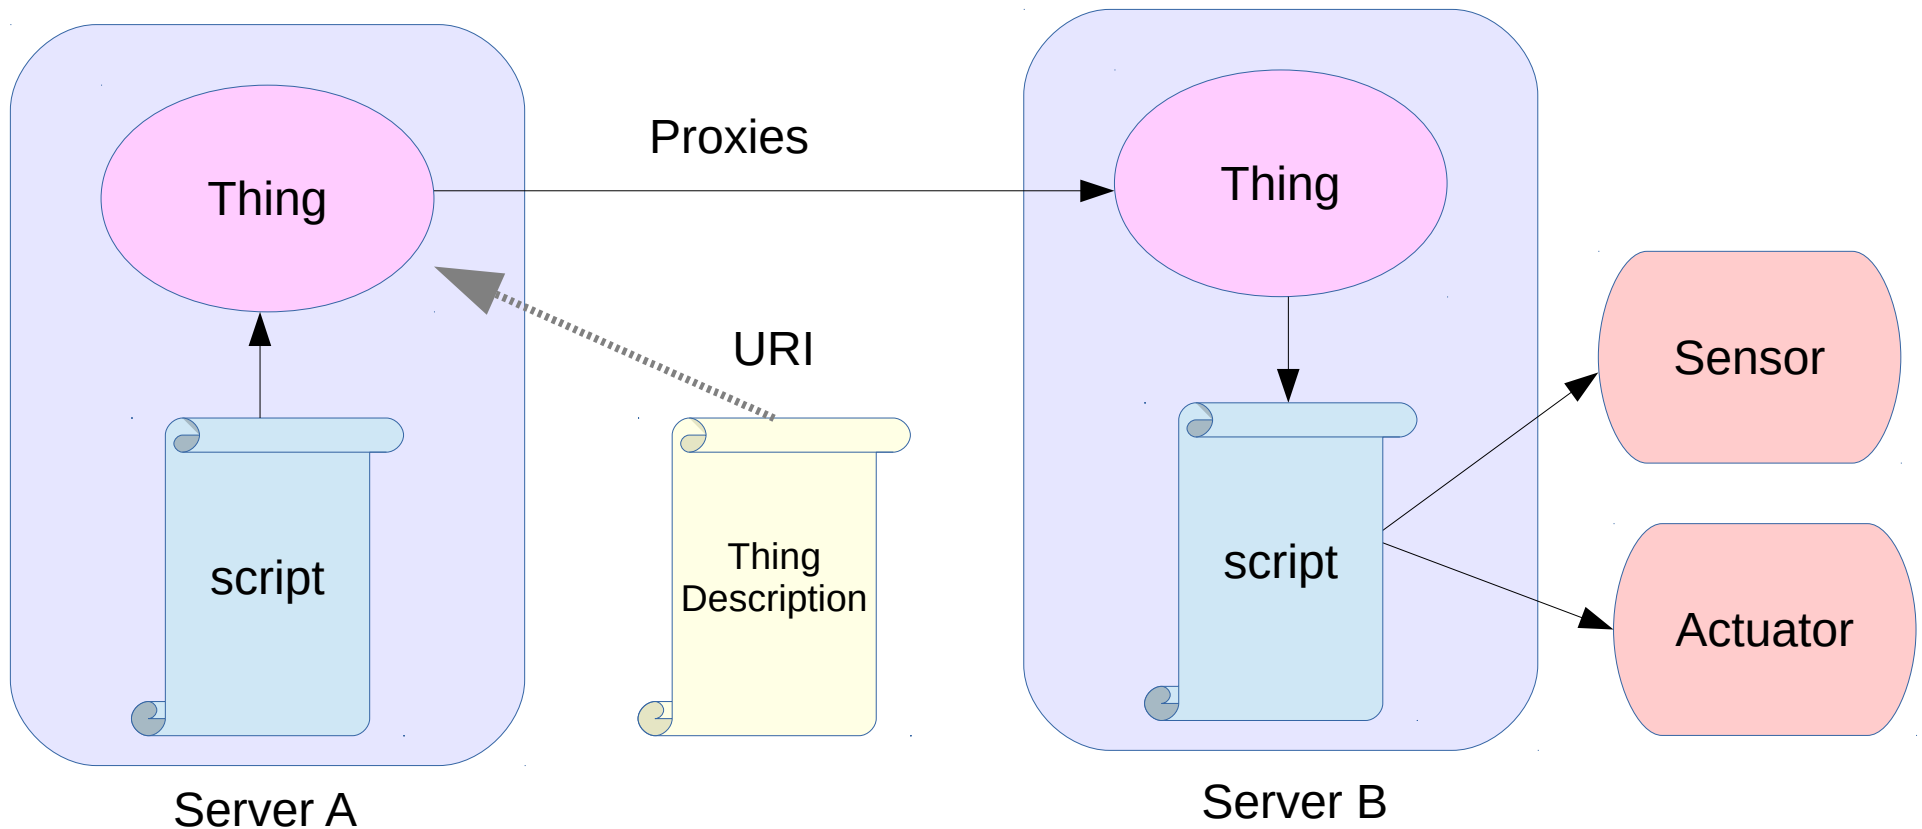
Distributed Web of Things

- Virtual representations of physical or abstract entities for use by application scripts
 - Each thing has a URI for a description which is used by the server to create a software object for applications scripts to interact with in terms of the thing's properties, actions, events and metadata



Distributed Web of Things

- Thing descriptions can be used to create proxies for a thing, allowing scripts to interact with a local proxy standing for a remote entity
 - Web page scripts in browser can create proxies for things on servers



Abstraction Layers

Application	Scripts that define thing behaviour in terms of their properties, actions and events, using APIs for control of sensor and actuator hardware
Things	Software objects that hold their state Abstract thing to thing messages Semantics and Metadata, Data models and Data
Transfer	Bindings of abstract messages to mechanisms provided by each protocol, including choice of communication pattern, e.g. pull, push, pub-sub, peer to peer, ...
Transport	REST based protocols, e.g. HTTP, CoAP Pub-Sub protocols, e.g. MQTT, XMPP Others, potentially including non IP transports
Network	Underlying communication technology with support for exchange of simple messages (packets)

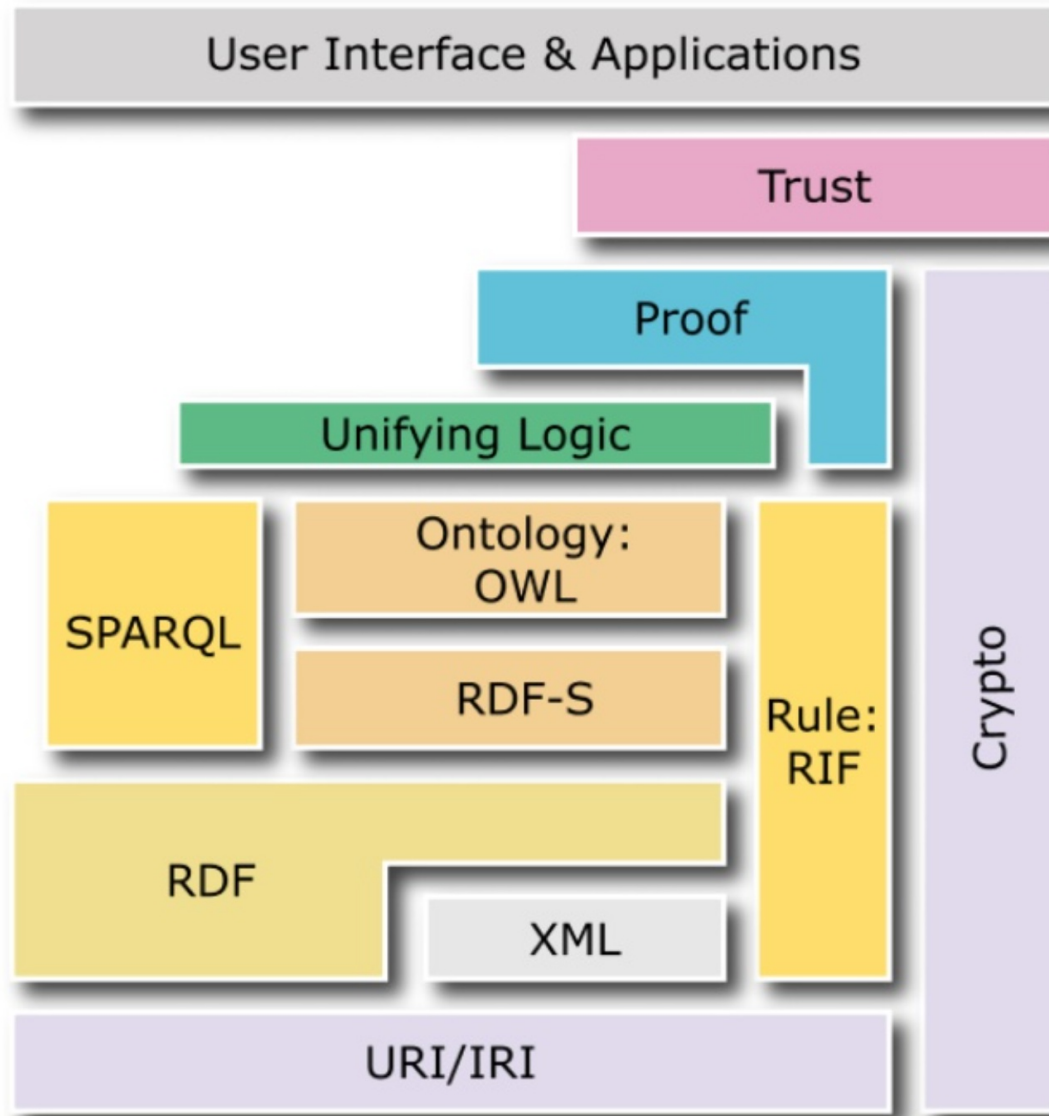
Semantics – a very quick recap

- Semantics is the study of meaning
 - The relationship between words and what they stand for
- Some things we can talk about
 - **People**, e.g. Robert Schuman
 - **Places**, e.g. Brussels
 - **Events**, e.g. the formation of the European Coal and Steel Community (ECSC)
 - **Dates**, e.g. 9 May 1950
 - **Documents**, e.g. the Schuman Declaration
- In other words, a mix of physical and abstract entities

Relationships Between Entities

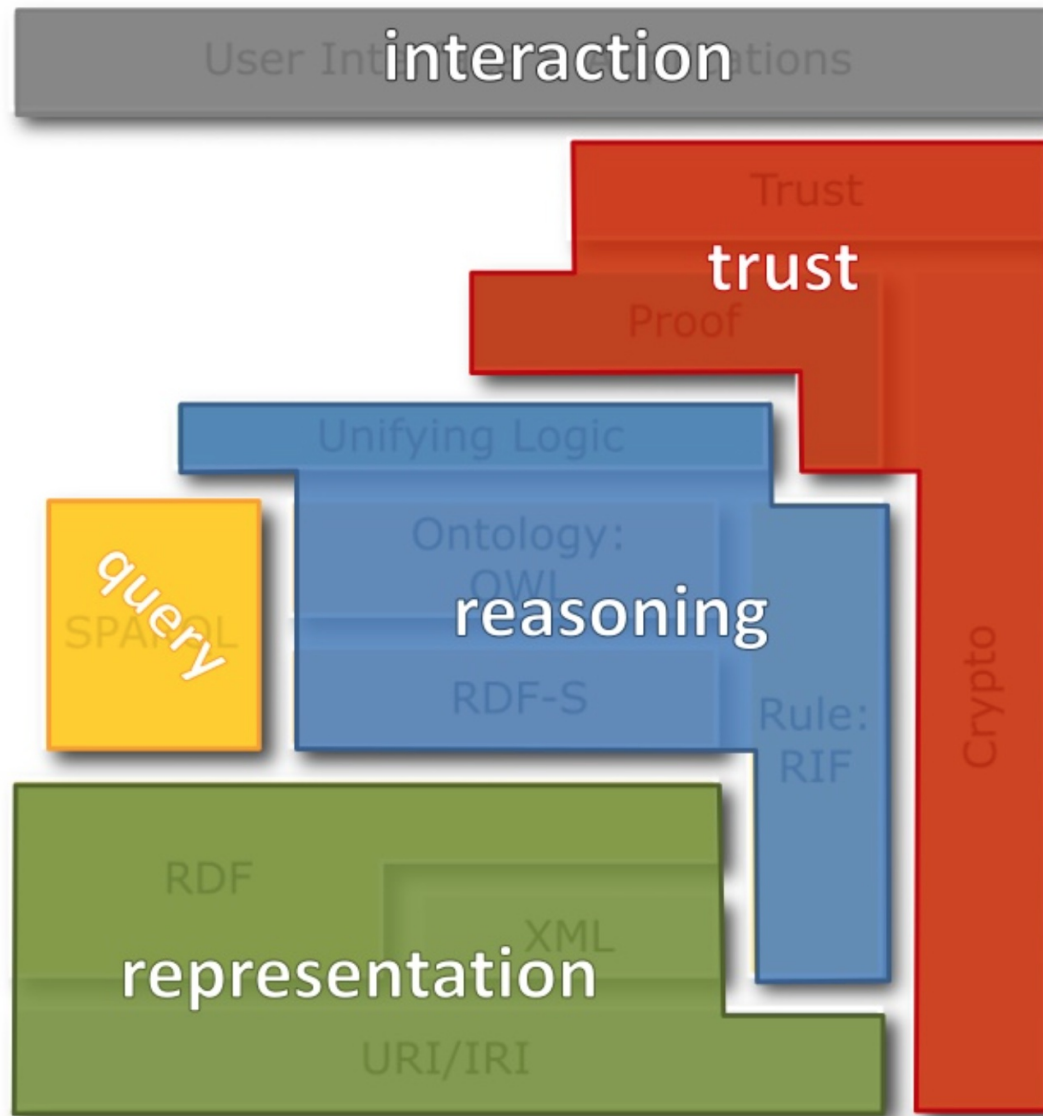
- Named relationships between entities
 - **Brussels** is the **capital city** of **Belgium**
 - Subject: “Brussels”
 - Predicate “capital city”
 - Object: “Belgium”
- W3C's **Resource Description Framework**
 - Subject, Predicate and Object as Web addresses (URLs)
 - These URLs act as globally unique identifiers
 - The URLs can be dereferenced to obtain further information (hence the term “Linked Data”)
 - RDF has many serialisations, e.g. XML, Turtle, JSON-LD

W3C Semantic Web Stack




With thanks to Fabien Gandon

W3C Semantic Web Stack



A Growing Cloud of Linked Data



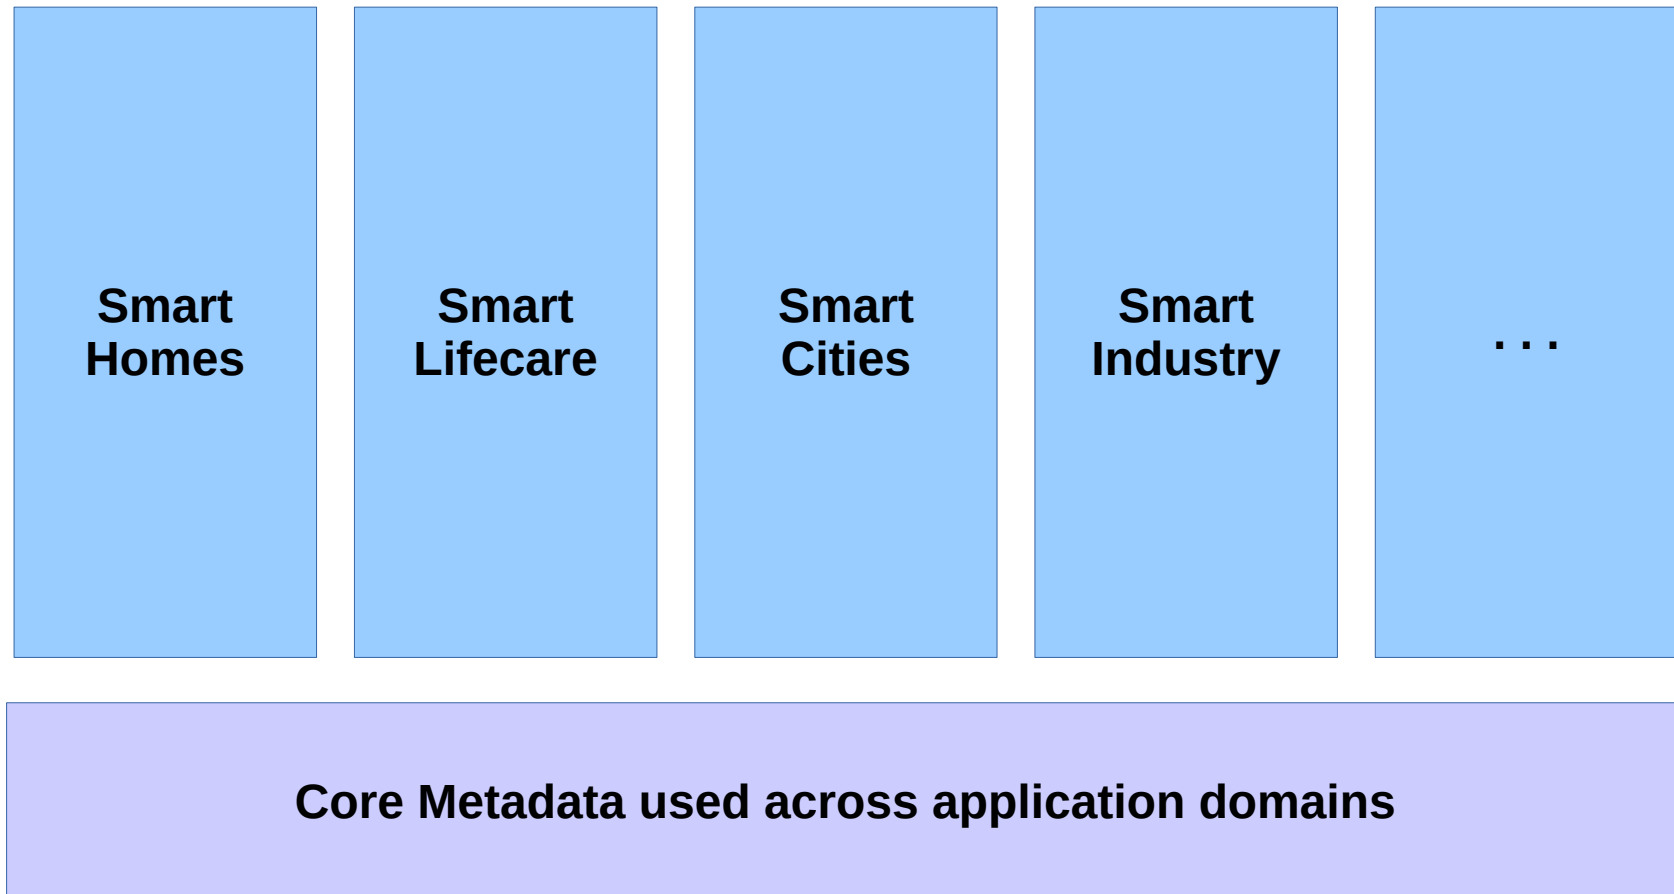
- Publications
- Life Sciences
- Cross-Domain
- Social Networking
- Geographic
- Government
- Media
- User-Generated Content
- Linguistics

Linked Datasets as of August 2014

Semantics and the IoT

- What is the relevance to the Internet of Things?
 - Shared vocabularies for entities and their relationships
 - Describing the software objects that stand for “things”
 - Verifying that a data source and sink are compatible and have the same semantics
 - Floating point number representing a temperature value expressed in Kelvins
 - When searching for services with a given semantics
 - Show me all temperature sensors in a radius of 100m
 - To facilitate the design of service compositions
 - To enable simulation prior to deploying changes to cyber-physical systems
 - To enable fault diagnosis based upon causal models


Horizontal and Vertical Metadata



Industry specific groups are in best position to define metadata for each vertical

W3C view of Horizontal Metadata

Core metadata applicable across application domains



- **Thing descriptions**
 - Links to thing semantics
 - Data models & relationships between things
 - Dependencies and version management
 - Discovery and provisioning
 - Bindings to APIs and protocols
- **Security related metadata**
 - Security practices
 - Mutual authentication
 - Access control
 - Terms & conditions
 - Relationship to "Liability"
 - Payments
 - Trust and Identity Verification
 - Privacy and Provenance
 - Safety, Compliance and Resilience
- **Communication related metadata**
 - Protocols and ports
 - Data formats & encodings
 - Multiplexing and buffering of data
 - Efficient use of protocols
 - Devices which sleep most of the time

Web of Things Topologies

- The Web of Things lends itself to different topologies
 - **Peer to Peer**
 - Devices talk directly to one another
 - Each device can host a mix of things and proxies for things on other devices
 - **Peer to Peer via Cloud**
 - Using a message routing network
 - Using WebRTC data channel
 - **Star** – Hub as controller for cluster of devices
 - The hub has proxies for the things on each of the devices
 - **Device to Cloud**
 - Device registers things on Cloud-based server
 - **Star to Cloud**
 - Hub acts as gateway between devices and the cloud

Intelligence in Depth

- Abstraction layers for **sensing**
 - Progressive stages of interpretation
 - Combining sensor data with other sources of information
 - Inferred events
 - Machine learning
 - Monitoring to check all is well
 - Reducing the burden on cloud based systems
- Abstraction layers for **actuation**
 - Progressively map high level intent to low level actuation
 - Synchronisation across clusters of devices
- Abstraction layers for **control**
 - Control links sensing to actuation
 - Implementing control at multiple levels of abstraction

Web of Things at W3C

W3C Web of Things Interest Group: <http://www.w3.org/WoT/IG/>

- W3C Workshop in Berlin in June 2014
- Launch of Web of Things Interest Group in 2015
 - Task forces
 - Thing descriptions
 - APIs and protocols
 - Discovery and provisioning
 - Security, privacy and resilience
 - Communications and collaboration
 - Emphasis on implementation experience
 - Demos and plugfests
 - Face to face meetings
 - Past: Munich, Sunnyvale, Sapporo
 - Joint meetings with IRTF Thing to Thing Research Group
 - Future: Jan '16 Nice, France, April '16 Cambridge MA, July '16 Asia, September '16 Lisbon, Portugal

W3C Web of Things Interest Group



Plans for Launching Web of Things Working Group

- The Interest Group is working on
 - Use cases, requirements, technology landscape and plans for launching working groups
 - W3C Interest Groups prepare the ground for standards but don't develop standards
 - W3C Working Groups are chartered to develop standards (W3C Recommendations)
- We're collecting ideas including
 - Horizontal metadata vocabularies
 - things, security, communications
 - Serialisations of metadata, e.g. as JSON-LD
 - APIs and bindings to specific protocols & platforms
- We expect to launch the WoT WG in 2016

Web of Things & M2M

- The Web of Things as a platform of platforms
 - Focus on simplifying application layer
 - Delegate details to servers when it comes to the protocols, messages and communication patterns
 - This is possible through rich metadata
 - W3C's expertise with the Semantic Web
- M2M should be integrated as one of the platforms
 - RESTful messages for HTTP, CoAP and MQTT
 - But this requires a way to map M2M's resources to the application layer semantics with properties, actions and events
 - ETSI/oneM2M should work with W3C on this mapping

Web of Things for Developers

- The Web of Things is designed to appeal to developers by reducing the cost and complexity for creating services
 - Clean separation of abstraction layers
- A uniform approach to services that scales across devices, platforms and domains
- This will pave the way to an open market for services on the scale of the Web
 - Increased market size for your solutions
- Help W3C to address the challenges and lead the Web of Things to its full potential!

Discussion?



Many Protocols

- Internet Protocols

- HTTP*
- Web Sockets
- CoAP*
- MQTT
- XMPP
- AMQP

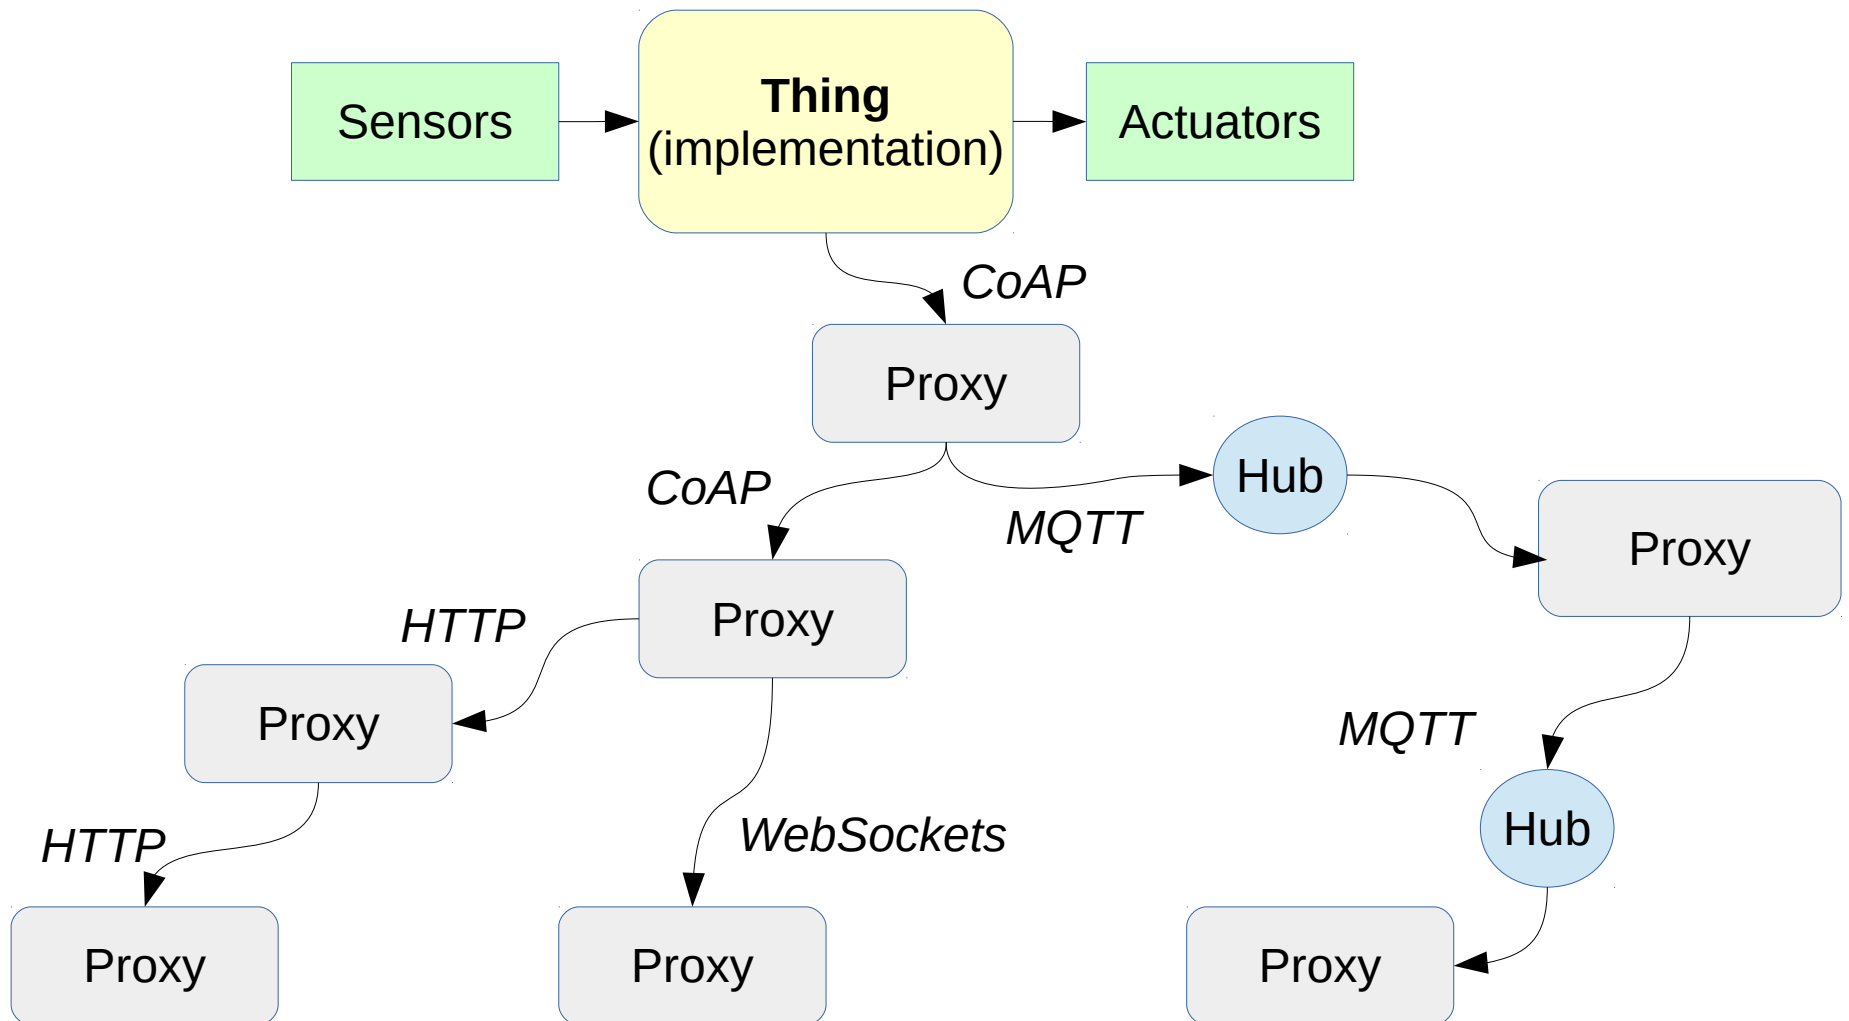
- IoT protocols

- CoAP over 6LoPAN
 - IPv6 over 802.15.4
 - IPv4 & IPv6 over WiFi
- MQTT-SN
- Bluetooth Smart (BLE)
- ZigBee
- KNX
- EchoNet
- ETSI LTN, Weightless, LoRaWAN, SIGFOX UNB, ...
- *and a great many more*

* Commonly used with REST

Metadata as basis for decoupling services from protocols

Using a heterogeneous mix of protocols



Driving convergence on standards

- Groups like IETF, ETSI, oneM2M and OIC focus on RESTful interfaces
 - The IoT is modelled in terms of a hierarchy of resources
 - Applications need to know about the communication patterns
 - Weak semantics
- Compose focuses on streams and data models
- W3C addresses the application layer
 - Generalises Compose approach to platform of platforms
 - Decoupling applications from the protocols and communication patterns
 - Strong focus on strong semantics and rich metadata
- On-going discussions on convergence
 - Collaboration agreements with IIC, oneM2M, OIC, IoTSEF, Industry 4.0, ...
 - Joint W3C & IRTF Thing to Thing Research Group meetings (Prague, Yokohama)
 - W3C invited presentation at OIC conference, 5 November 2015
 - W3C invited to present at ETSI M2M workshop, 9 December 2015

Dealing with Challenges to the IoT?

- The relationship between different industry alliances and standards development organisations. Why are there so many, and what is needed to drive convergence and reduce the fragmentation that is holding back the potential for the IoT?
- The difficulty of getting people from different communities and backgrounds to work effectively together. People essentially see the world through the prism of their own experience. This explains the hammer and nail phenomenon, where if you are emotionally attached to a hammer, everything appears to look like a nail. Such communities tend to ignore problems that don't fit nicely into their world view.
- Large companies quite naturally want to create and control their own ecosystems, and may be lukewarm when it comes to support for work on vendor neutral open standards.
- How to counter the risk that the network effect will enable a few companies to dominate globally?
- How to counter the risk that national governments will use the IoT to monitor every aspect of our lives and repress minorities and politically disadvantaged members of society?
- How to counter the risk that companies collecting vast amounts of personal data will abuse this through an imbalance of power between corporations and individuals?
- How to counter the risk of cyber attacks as society becomes increasingly dependent on the IoT?