www.pipelinepub.com
Volume 18, Issue 2

The Technology Roadmap to Enable IoT-Data Conversations

By: Ken Figueredo, Joachim Koss

When attending a cocktail party or similar social event, a person can expect to meet old friends. They might also make completely new acquaintances via friends of friends or the host's introductions. Some of these new acquaintances might be people from out of town. Others might speak in a different language, requiring a translator or third language for communication.

There is no reason why this human-centered scenario should not apply to connected devices, sensors and IoT applications. In the case of consumer IoT components, the equivalent of the party gathering might be a home. Other environments might be office buildings, cities, factories, and even transportation networks. In each of these places, the growing proliferation of IoT applications and devices creates new opportunities for cross-silo applications. Like the party host, there is also the potential to economize on hardware and communications costs by using one IoT sensor to serve data to multiple applications.

Two factors are involved in enabling 'conversations' between IoT devices and applications. The first is a common framework that allows different IoT devices and applications to communicate with one another. This is generally one of the functions of an IoT platform. The second requirement is a shared set of rules that allows devices, sensors, and applications to exchange IoT data. This comes under the discipline of semantic interoperability.

Interoperable framework for IoT platforms

In 2012, a group of leading standardization bodies launched the oneM2M initiative to establish a standard for end-to-end and interoperable IoT systems. These bodies wanted to avoid regional fragmentation and promote a global IoT market, reflecting the foundations and successes of the mobile telecommunications industry.

The oneM2M standard addresses situations where an organization can deploy an IoT solution, using components from different suppliers, and then add other solutions over time. The use of an IoT middleware capability between sources and consumers of IoT data aims to mask complexities along the IoT technology stack. In effect, the oneM2M standard defines a three-layer, horizontal architecture that is reusable across different application sectors. The lower layer corresponds to devices and communications technologies. The upper layer corresponds to applications that process data from IoT devices and sensors for decision-making and control interventions. oneM2M's technical specifications define a set of common service elements that reside in the middleware layer.

One way to think about these common services is as a set of technical capabilities that application developers can use to design and deploy IoT systems. Many of these are common to all IoT applications. Take the example of the 'registration' service. A developer could use this to establish the authorization and authentication relationships between different device, gateway, platform, and application entities in an IoT system.

Another common service function is security. Here, oneM2M defines a common approach for the handling of sensitive data, security administration, establishment of security associations, access control (identification, authentication, and authorization) and identity management. One of the benefits of oneM2M's approach is that developers can use and reuse the set of common services to build applications for public safety, smart cities, intelligent transport, and other sectors. In addition, oneM2M's standardization framework allows new common services to be added to the toolkit as new requirements emerge over time.

Enhancing IoT systems for semantic interoperability

Having overcome issues of basic connectivity and network interoperability through the horizontal architecture and common services, syntactic interoperability becomes the next challenge. This relies on the use of information models to represent IoT devices, using static information based on a pre-defined syntax. In the case of a heat sensor, for example, the pre-defined syntax might specify that data is sent as a floating point, temperature reading.

Once requirements about information exchanges becomes more complex, as is the case with systems from different domains, static information is no longer sufficient. Now, there is a need to base the exchange of information on its meaning, independent of underlying protocols.

Returning to the heat sensor example, semantic attributes might indicate that the reading comes from an indoor sensor that is situated on the second floor in a specified commercial building.

The next complication arises when data sources and data-requesting entities use different ontologies because they are developed by different companies not using a common ontology. When devices or a service application want to read and understand data, semantic interoperability is needed. It combines the ability to establish a shared meaning of the data exchanged as well as the technologies to interpret communication interfaces.

Advances in semantic interoperability

Existing technical specifications in the oneM2M standard contain the basic functionality for semantic interoperability. An application developer can invoke this as a service to allow an application to discover, connect and collect time-series data from an IoT sensor.

oneM2M's current discovery capabilities work properly only if the discovery is well-scoped and designed (for example, an attribute field indicates that a set of lights are in a house) or if a search is related to specific known sources of information (for instance, searching for the values of a known set of application data records). When oneM2M is used to discover wider sets of data or unknown sets of data, the functionality is typically integrated by *ad hoc* applications to expand the oneM2M functionality. This is not optimal for interworking and interoperability when this core function is implemented in assorted flavors.

One way to improve matters involves a more dynamic approach to semantic interoperability. This is the topic of an ongoing project in European standards organization ETSI. It aims to enable an easy and efficient discovery of information and a proper interworking with external sources and consumers of information. This might apply to discovery in a distributed data base belonging to a smart city or an intelligent factory. It might also involve a direct search for information in the oneM2M system for big data purposes.

Returning to the cocktail party scenario, the new capabilities correspond to a person having the ability to find new people, who share similar interests, and to have a meaningful conversation over a common topic. Enhancements identified through the ETSI project will be introduced as new functionalities in Release 5 of the oneM2M standard.

Future trends and the residual value legacy systems

In the context of IoT systems, interoperability makes it possible to build applications that work across domains. The same principle makes it possible to develop applications using components that are supplied by different vendors. Data and system interoperability will become increasingly

important in IoT systems as greater quantities of data are generated and shared across users and IoT platforms. Semantic interoperability enables new commercial opportunities in different domains such as healthcare, smart grid, smart metering, intelligent transport systems, industrial automation systems, and smart cities. Each of these examples relies on the collection, processing and sharing of data across organizational and operational boundaries.

A current work item in oneM2M on information models for the rail sector illustrates these ideas in a real-world setting. Consider a safety-related use case to alert train drivers in emergency situations based on capturing data from various kinds of railroad crossing gates along a railway line.

Another kind of gate and use case concerns ticketing and flow-control systems based on smart gates in train stations. In this case, some of the attributes in the gate's information model deal with the open or closed status of a gate as well as data related to barcode and e-ticket scanning. Other attributes might deal with passenger instructions that might be sent to digital displays in the electronic gates. oneM2M defines a device template that provides a standardized way for developers to interact with IoT devices. In addition to covering a basic set of attributes, the template has the flexibility to include other attributes.

To date, Release 3 of the oneM2M standard established a library of 84 model classes for home domain devices. Additions to Release 4 go beyond the home domain to other verticals with 119 module classes and 19 common device models spanning home, health, city, vehicular and railway domains, as well as two sub-device models. This library makes it possible to form direct links between smart home and retail or manufacturing domains, for example, thereby enabling new, cross-domain services.

In addition to the future-oriented roadmap to enable new IoT capabilities, the oneM2M standard also addresses applications that depend on legacy technologies. A case in point is OPC-UA, which is a widely used industrial automation protocol for switches and programmable logic controllers (PLCs). oneM2M members defined an Interworking Proxy Entity (IPE) that provides syntactic interoperability between OPC-UA and oneM2M domains. It lays the foundation to support interworking between the OPC-UA information model and the oneM2M information model.

Building on this foundation, developers can use oneM2M's Semantic Reasoning capability to interact with legacy devices. This is an example of oneM2M adding a layer of value on top of OPC-UA's data exchange functionality. As IoT systems become more common, this example illustrates another long-term industry requirement to capture the value of legacy systems through new interoperability capabilities.