www.pipelinepub.com
Volume 18, Issue 9

Making IoT Intelligent

By: Ken Figueredo

In team sports, each player has a particular role. Team effectiveness depends on individuals performing as a system. This means all members work together as a tight-knit unit. A lack of cohesion leads to inferior performance, which might result from players not being set up with a 'game plan' or individuals being unable to adapt their play intelligently.

Internet of Things (IoT) systems are not dissimilar to sporting teams. They involve combinations of connected devices and sensors. Each has a role, either to act on command or to supply data to applications where decision-making occurs. If a hardware component from one vendor breaks down, it might be substituted with a comparable item, possibly from a different vendor. If the second vendor uses a proprietary interface or data communications protocol, this introduces friction into the overall IoT system. It is a bit like bringing a new player into a team.

Ideally, it would be useful to work with intelligent components when building end-to-end IoT systems. That way, a new component could introduce itself and announce its capabilities to the system. In IoT terminology, this is known as a ‘discovery’ process. It is linked to the ‘registration’ process, which deals with credentials management. This allows components to identify and authenticate themselves in IoT systems. These are like the social conventions that all players use when joining a team. In the IoT context, ‘discovery’ and ‘registration’ are examples of common and reusable functions that add intelligence to IoT components.

Treat IoT as a system

In designing and working with IoT systems, it is useful to think about how different components fit into a three-layer stack. The lower layer encompasses an array of IoT devices, sensors, and network communications technologies. The upper layer contains software applications that consume data from IoT devices and sensors. Their role is to spot patterns, make decisions and issue control instructions. The middle layer functions as a technology-abstraction capability. It enables interactions between upper- and lower-layer components. This is where common functions reside, such as ‘discovery’ and ‘registration.’

In technical terms, the middle layer is associated with IoT platforms. Middleware capabilities are offered as callable services to different components in an IoT system. This means all connected devices, sensors and applications use a common language and protocols to interact with one another. It is a bit like the social conventions and tactics used by players on a team.

Real-world IoT complexities

In real-world deployments, the three-layer framework still faces the challenge of a fragmented technology landscape. Take the example of a greenfield deployment with no legacy encumbrances. Here, a customer has one of two implementation approaches. One is to go for a full solution, from the devices to the service application, drawing on a small and pre-integrated set of technology options offered by a solution supplier. The alternative is to go for different combinations of devices and applications from the open market. This second approach requires some extra technical 'glue' to integrate different subsystems and any incompatibilities that might exist between them. Here, the design challenge is to find a way through a jungle of technical choices and engineer a system with integration solutions for each interface point.

'Brownfield' deployments, of course, introduce an additional layer of complexity. Take the case of an industrial processing plant in which the operator wishes to add new IoT capabilities to an existing production facility. This involves an integration of new and legacy technologies to capitalize on an asset's remaining life and potentially unique functionality. One approach is to integrate the two sub-domains via a customized gateway function. This might work in small systems. For larger systems, it becomes costly because of technology permutations and the staffing expertise required for future maintenance activities.

IoT platforms aim to alleviate these issues. However, as the IoT market has grown, so has the number of IoT platforms. There are as many as 1,600 platforms to choose from. Any enterprise on the path to adopting IoT strategically is playing roulette when choosing which platform to adopt.

Another industry approach involves the use of application programming interfaces (APIs) to enable interoperability. With users embracing this technology, there has been considerable growth in the market for API libraries and service providers. There is a lesson to be learned, however, from the market for consumer data and the growing concentration of market power of large platform providers. There, the use of open APIs is viewed as a means of enabling data interoperability. What would be better, however, is an open standards approach.

Open standard for IoT systems

In 2012, a group of national and regional standardization bodies across the world launched the oneM2M initiative. Participants from the Americas, Asia, and Europe set out to establish a standard for interoperable IoT systems, avoiding regional fragmentation and promoting a global IoT market. The standardization process began with an analysis of common requirements across different industry verticals. It took a systems approach and planned for a roadmap that could accommodate future use cases and requirements.

oneM2M's roadmap began with its Release 1 technical specifications in 2015. These provide basic connectivity and security features that any connected-device or application in an IoT system can call upon. These were the first in the portfolio of common service functions. Since then, over two hundred organizations across the world added to the standard, addressing new IoT requirements and integrating emerging technologies in areas related to security, privacy, and AI.

Being an open standard, oneM2M provides the 'glue' to deploy components from different suppliers in IoT solutions. They can add other components to grow their solutions over time. They can also connect separate systems to link application silos as well as departmental and organizational boundaries. This might arise when linking congestion management, fleet transport and pollution monitoring systems for intelligent transport and smart city applications, for example.

The standard continues to evolve and incorporate new requirements. oneM2M will soon issue Release 4 of the standard while working in parallel to plan Release 5 capabilities. These address the combination of artificial intelligence and machine learning techniques with IoT, advances in semantic interoperability and enhancements to support data protection regulations such as GDPR.

Roadmap for intelligent IoT

Connectivity and communications are present-day challenges in the IoT market. This is evident from satellite IoT to connect geographically distributed systems. Edge computing is the hot topic to connect highly localized computing situations.

In tomorrow's world, however, innovation in IoT systems will shift to making component technologies more intelligent. This is part of the journey to making better use of data. Examples of intelligent components include tiny sensors with limited processor capabilities and battery life. These might be embedded in the fabric of buildings and bridges to enable safety monitoring and predictive maintenance applications. To prolong their service life, these devices need to operate intelligently. This means knowing when to transmit meaningful data, principally when sensor readings change. Another form of intelligence is to know when to enter a deep-sleep mode for energy conservation purposes.

Future requirements for IoT data also go beyond better decision-making. There is much greater value from sharing data with different users and even external stakeholders. Consistency in recording and communicating data are examples of new requirements. The path to these opportunities involves technical standards for information models and digital twins of connected devices and sensors. It is also important to recognize that libraries and streams of real-time data will belong to different users and be valued differently. This will lead to new requirements to share data intelligently using control policies and to enable monetization through selective sharing and licensing mechanisms.

As services based on IoT systems become mission-critical and add to the convenience of everyday living, users will develop higher expectations about quality, consistency, and transparency. This translates into a set of requirements for data provenance and the credentials of IoT data sources

and sensors. To address these, there is already growing interest in explainable and trustworthy AI for IoT systems.

These more sophisticated use cases rely on extra intelligence in IoT components, something that product managers need to factor into their plans. The use cases also depend on greater intelligence at a systems level. This ensures different components can work together with greater fluidity, much like a high-performing team.