

IoT Needs Open Discovery Scheme

By Francis DaCosta, 08.26.14 3

Share Post

The Internet of Things will require an open-source self-classification scheme if it is to scale to its full potential of trillions of nodes.

A great many people fail to understand the fundamental changes that will come with the Internet of Things. Instead, they expect traditional networking protocols such as TCP/IP and host-server architectures to carry on. After all, doesn't IPv6 make possible trillions of addresses? And isn't that the only problem that was left to solve in the Internet of Things?

That might be true enough if the only devices on the IoT were the PCs and smartphones familiar to us today. But that's not the future. Instead, there will be myriad purpose-built machines, sensors, appliances, and actuators pumping out countless streams of data. This rise of the machines creates the need for an open-source solution in three ways.

First, most of these relatively simple devices will be cheap -- less than a few dollars per device. They simply can't absorb the costs of traditional networking.

Second, they will number in the trillions. No current manmade scheme can digest data from that many devices.

Third, their manufacture and sale will be completely decentralized. Millions of manufacturers sprinkled throughout the world will lack the manpower and expertise to interact with centralized databases or standards bodies, so there is no hope of distributing media access control IDs to these makers.

The answer to gaining useful information from this impending fire hose of IoT data is simplicity itself. It requires the use of publish/subscribe architectures in IoT servers.

This is key because the most interesting data streams may be completely unknown to the organizations that need them. By searching for published streams of a particular type, streams from a defined area, or streams that vary in an interesting way, servers may build up information neighborhoods to extract and integrate knowledge from the IoT.

To discover which streams are of interest, we need to turn to nature, the one system that has already solved problems on a scale this massive. The fundamental concept in nature's systems is self-classification. By identifying themselves as a particular type of sensor, applicant, or actuator, simple IoT devices blindly bleating out values and readings can provide the necessary information to be incorporated into a higher-level picture of the world through the subscriptions of servers.

Some of these data streams might be known structures, such as those of an OEM that define the signatures of its own devices. But many more -- and perhaps the more interesting ones -- will be built up from previously unknown devices as servers subscribe to pertinent data streams.

Everyone and no one will manage this self-classification scheme. By making these self-classifying data structures open-source, they may readily be adopted by anyone who wishes to participate in the IoT. As I outlined in [my book](#), this open-source taxonomy is combined with the capacity of private extensions to permit some parts of the data stream to be available to all, while other parts are restricted to those with the correct key.

Self-classification is the fundamental requirement for publish/subscribe architectures to work, and publish/subscribe will be the only way to make sense out of the emerging Internet of Things. But that self-classification must be open-source if it is to scale to the trillions of devices and millions of makers that are in our very near future.

-- [Francis DaCosta](#) is an engineer, architect, and author focusing on scalable architectures for the emerging Internet of Things.