Connecting Fog and Cloud Computing

I ATTENDED THE INTERNET OF THINGS WORLD 2016 EVENT IN SANTA CLARA, CALIF., last year and served as the track chairman for, you guessed it, cloud and IoT. The feeling I got throughout the event was one of confusion: IoT seems to be so systemic, yet is difficult to define. As one presenter put it, "It's like plastic. It's going to be a part of everything."

According to Research Nester, "The Global Internet of Things (IoT) market reached USD 598.2 billion in 2015 and the market is expected to reach USD 724.2 billion by 2023. Further, the market is projected to register a CAGR of 13.2 percent during the forecast period 2016-2023 globally."[1]

But you don't have to tell us. Everything in our lives from the cars we drive, to the thermostats on the wall, to our refrigerators, literally, has a mind of their own these days. So much so, that we have ridiculous examples, everything from connected egg trays to Bluetooth-enabled toilets.

DAVID S. LINTHICUM

Cloud Technology Partners
david.linthicum@cloudtp.com

With all that said, we do have challenges to consider. The challenges include how we can get data processed from so many external devices. According Cisco, cloud traffic is likely to rise 3.7-fold by 2020, increasing 3.9 zettabytes (ZB) per year in 2015 (the latest full year for which data is available) to 14.1 ZB per year by 2020.[2]

Moreover, big data-associated Internet of Things devices are a large cause of this growth. Indeed, by 2020, database, analytics and IoT workloads will account for 22 percent of total business workloads, compared to 20 percent in 2015. The total volume of data generated by IoT will reach 600 ZB per year by 2020, which is 275 times higher than projected traffic going from data centers to end users/devices (2.2 ZB); 39 times higher than total projected data center traffic (15.3 ZB), according to the same Cisco report.

Thus, we have the perfect storm of the use of cloud computing, and the growth of IoT. IoT is about processing data that comes from devices in some way that's meaningful, and cloud computing is about leveraging data from centralized computing and storage. Growth rates of both can easily become unmanageable. We have some problems to solve.

Define the Problem, and the Solution

In the context of the Internet of Things, the trouble with the cloud is that data needs to be sent back from the sensors gathering info, such as a Nest thermostat or a Fitbit wristband, to a database in a remote public cloud. The time that it takes for the data to be transferred from the device or sensor to the remote public cloud, that is the latency, is often too great to meet the requirements of the IoT system.

We need to do something different, and we can start by doing IoT applications at the cloud's edge. This means that we avoid sending all the data from sensors and devices back to the cloud, but instead build data and applications on the edge of the network that can handle most of the data gathering and processing.

Recently, I published an article in *Computer* about Edge computing[3]. In that article I defined an architecture called Responsive Data Architecture, in which I mentioned that IoT brings this issue of moving some computing to the edge again. For example, say there is a machine on a factory floor that

2325-6095/17/$33.00 © 2017 IEEE

analyzes the quality of an auto part that it makes. If the part is not up to quality, as determined by an optical scanner, then it is automatically rejected.

While this keeps a human from looking at the part, and thus slowing down the process, it also takes a great deal of time to transmit the data and image back to the centralized database and compute engine, where a determination is made as to the success of the manufacturing process, and then communicated back to the machine.

The cloud complicates this process even more. We're focused on centralized computing, thus there will be latency. Now, instead of sending the data back to the data center on the other side of the factory, we send it to a remote cloud server that can be thousands of miles away. To make things worse, we send it over the open Internet. However, considering the amount of processing that needs to occur, the cloud is typically more efficient.

So what do we do? How do we solve the problem? We already know that computing at the edge pushes most of the data processes out to the edge of the network, close to the source. Then it's a matter of dividing the processing between data and processing at the edge, versus data and processing in the centralized system, meaning a public cloud such as Amazon Web Services.

The concept is to process the data that needs to quickly return to the device. In this case, the pass/fail data that indicates the success or failure of the physical manufacturing of the auto part. However, the data should also be centrally stored, and, ultimately, all of the data is sent back to the centralized system, cloud or not, for permanent storage and for future processing.

The benefit is better performance and efficiency. IoT applications need to react almost instantly to the data generated by a sensor or device, such as stopping a train, if sensors have reported problems with the track switch a few miles ahead, or shutting down an industrial machine that is about to overheat and explode. There are hundreds of use cases where reaction time is the key value of the IoT system.

Of course, we have to give this a name. Cisco Systems has tried to brand it fog computing and set up the OpenFog Consortium to promote its view. Whatever it ends up being called and defined, the key is to reduce latency for response-critical applications by moving the data transfer and processing to the edge of the cloud, closer to the IoT device.

I've been involved in dozens of systems where the data and applications were placed near the source, yet still working with centralized data and applications. While it's a bit tricky, it's not that hard to do. So, what value does OpenFog bring?

There are a few benefits that I see, including:

- A standard architecture and enabling technology that allows you to approach edge computing in a simple but consistent way.

> Computing at the edge of the network is, of course, nothing new – we've been doing it for years to solve the same issue with other kinds of computing.

- The ability to provide a good product development framework that network devices and software builders can follow, as well as influence.
- The ability to deal with security in a consistent way. Last year DDOS attacks took over devices, not computers, and now that everything is smart and has an OS, this will be a fact of life going forward.

OpenFog recently published a reference architecture that covers pretty much everything from Security to Programmability (see Figure). If this looks like it's been designed by committee, it's because it has. I did not find it useful.

Like other open standards, OpenFog gets things done through workgroups and committees. The danger here is that OpenFog could suffer from "too many cooks in the kitchen." The lack of interest in many standards came about due to lack of speed. However, OpenFog does have a good list of member companies (see www.openfogconsortium.org/what-we-do/).

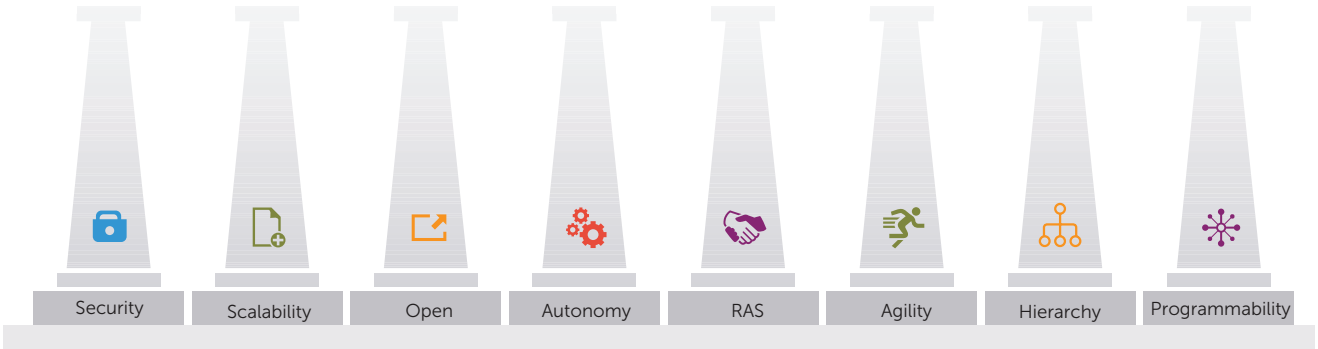

FIGURE 1. The OpenFog Reference Architecture is based on eight pillars.

What does this mean?

With the Internet of Things, the latency issue is more acute and more widespread than it is for other kinds of computing. That's why putting IoT at the edge of the cloud is such an important concept. Again, it's not that hard of a concept to carry out. Most distributed computing developers are very familiar with the concept of placing the processing as close to the source as you can.

No matter how speedy the networks get, latency will always be something that developers and admins will try to manage. While we can certainly toss new equipment at the problem, I've found that most performance issues need to be solved by changing the design, and not the infrastructure. This is the only way you can truly solve the problem.

So, the concept is sound, and OpenFog, and Fog computing, is attempting to formalize it, leading thought and promoting the notion of computing at the edge for cloud and non-cloud deployments. If the Cisco standard is successful, then OpenFog will have accomplished its objective.

That said, standards seem to fail, and this standard could be no exception. The fact of the matter is that they fail because so many of the member companies have their own agendas, which may not line up with the agendas of the other members. Thus, not much gets done, and the fruit of the standard dies on the vine. Fair warning, OpenFog.

If IoT and cloud are in your future (who does not have them in their future?), then you need to study this issue. This means reading my other article, Responsive Data Architecture (RDA), as well as understanding what OpenFog has to offer with an eye on what's realistic.

At their core, the value of all of these concepts is that we're considering alternatives to placing everything in the public cloud. Why? Because the public cloud does not make sense, in some cases. IoT will challenge us to think differently, and the use of edge computing, or fog computing, all combined with cloud computing, is the likely path that we will find ourselves upon.

References

1. Internet of Things (IoT) Market: Global Demand, Growth Analysis & Opportunity Outlook 2023. Published 1 February 2017, http://www.researchnester.com/reports/internet-of-things-iot-market-global-demand-growth-analysis-opportunity-outlook-2023/216
2. VNI Global Fixed and Mobile Internet Traffic Forecasts. 2017. http://www.cisco.com/c/en/us/solutions/service-provider/visual-networking-index-vni
3. David Linthicum, "Responsive Data Architecture for the Internet of Things", Computer, vol. 49, no. , pp. 72-75, Oct. 2016, doi:10.1109/MC.2016.302

DAVID S. LINTHICUM *is senior vice president of Cloud Technology Partners. He's also Gigaom's research analyst and frequently writes for InfoWorld on deep technology subjects. His research interests include complex distributed systems, including cloud computing, data integration, service-oriented architecture, Internet of Things, and big data systems. Contact him at david@davidlinthicum.com.*