10 Standards and Technology

Did I say this chapter was about standards and technology? Oops, sorry, I misspoke. This chapter is actually about making technologies work together for business benefits, which is IoT's sole raison d'être. Without seamless interoperability and integration there's no reason for IoT. All of those rosy projections—billions and billions in revenue created by millions upon millions of connected devices, communicating across vast numbers of networks, generating seemingly endless data for countless vertical applications—were based on the assumption that all of these elements, once they could communicate, would interoperate in a smart way. If they can't, who needs IoT at all? We'll just go back to the 20th century single-vendor-does-it-all custom solution models. This chapter will also take a closer look at a few key game-changing technology shifts we've been referencing throughout this book.

So this chapter is about standards, to the extent that only through universally accepted, efficient standards do all of those myriad things and parts and pieces and networks, both new and old, have any chance of communicating sufficiently up and down to exchange data, integrate, and

provide the right capabilities for the apps to generate business insights in open, interoperable, and industry-accepted ways. The standards efforts around IoT have been going on for many years now, beginning with the push to adopt and adapt to open networking technologies even before IoT came into vogue. Some of these standards have now reached the point of maturity and, yes, interoperability to enable the use-cases I've been citing in this book. Going forward, the industry is determined to evolve and improve standards within the next few years to keep pace with technology changes and enable more advanced use cases, applications, and value propositions.

The Case for Standards

With IoT, the technology situation is already much more versatile, complex, and fluid than the IT or OT scenarios with which you're now probably familiar. The number of legacy, proprietary, quasi-standard, and specialized technologies is simply mind boggling by itself. At the same time, the industry has begun to cull the traditional market structures of vendors out pushing end-to-end, proprietary, single-vendor solutions. If your company is one of those, become involved now in industry standards workgroups and learn how to thrive in an open world if you want to have a shot at thriving in the IoT-driven economy.

Let's add to that list the duplication of vendor ecosystems even within the same vertical. Just compare the tier-1 supplier ecosystem for German and U.S. automakers, and you'll start to see the picture. Further complicating the IoT standards are the sheer diversity of end devices, sensors, actuators, meters, controllers, appliances, and more that today have varied capabilities and proprietary device, management, and data interfaces and formats. And I haven't even started to count the number of embedded OSes, chip suppliers, and so on that further increase the complexity. This is one of the reasons why an entire industry of IoT platforms—the companies that connect and integrate with proprietary third-party end devices—has sprung up. They've resorted to creating their own abstraction layers and development environments through which application developers interact with these devices and the data they generate. While such an approach is needed today, it is inefficient and redundant. Yes, there

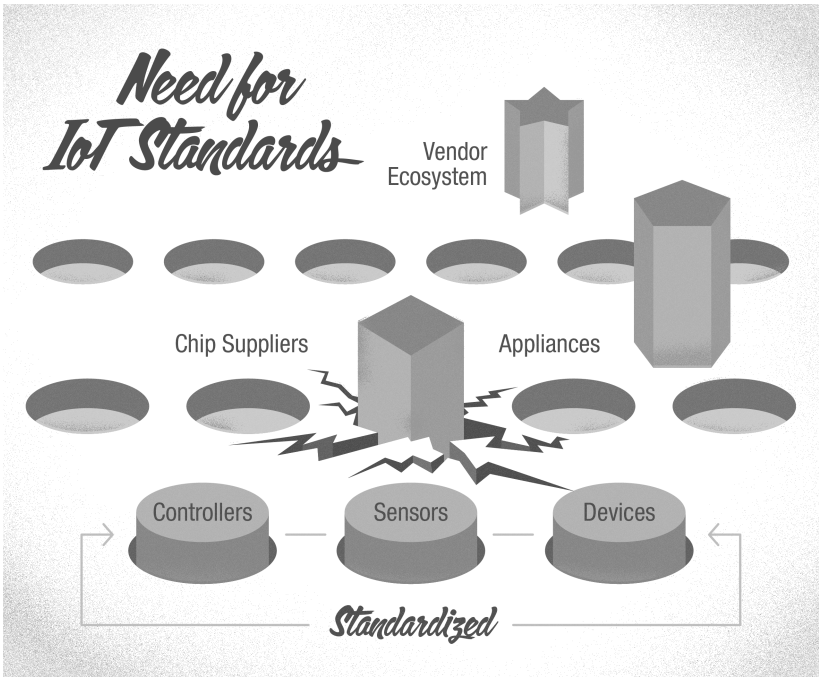

Figure 10.1 Need for IoT Standards

will always be a need for middleware to support legacy devices. However, getting the industry to standardize and adopt common data formats and APIs for new devices will be a big step forward (Figure 10.1). We urgently need to do just that.

Imagine how expensive such an approach is for each of these platform companies to undertake on its own. Even worse, you don't do it only once. Every time one or another of these devices or interfaces changes, you have to revisit it. That's why even the biggest players in the industry are willingly joining standards development teams. As much as each would love its own approach to become the accepted standard, it's too costly to develop and, most importantly, maintain by itself over time. The economics of doing this kind of standards work as an industry-wide collaborative effort are just too compelling to ignore.

But wait, there's more. As we've said throughout this book, we already have billions and billions of connected devices—including cars, buses, trains, office buildings, factories, oil rigs, homes, and entire cities.

Some are stationary, some mobile; some have IP addresses, others don't; some are always on, some intermittent; some are clustered together, others geographically dispersed. And that's just the beginning. The solution requirements vary vastly, as well. Some require devices to transfer megabytes of data every second, some just a few bits every few days; some want data to be analyzed in real time, some don't; some can be powered from the grid, some need to operate 20 years on a single battery. I hope you get the picture.[1] All of these variables are driving an interesting phenomenon. Unlike in the original Internet, we're actually seeing a proliferation of access or "last-mile" technologies. (Last-mile refers to the final leg from the network to the device.) No longer limited to Ethernet, Wi-Fi, and 3G/4G, IoT deployments today also include satellite, Bluetooth LE, low power wide area network (LPWAN) technologies such as LoRa, power line communication (PLC), and various wireless personal area networks (WPAN) such as Wi-SUN and ZigBee NAN, among many others. Which technology is best for each situation depends on several criteria, which we'll discuss later in this chapter. My first point here is to simply convey the complexity and vastness of the IoT world as it stands today at this early stage of maturity.

My other point is that, right now, you'll want to start the migration from proprietary and semi-proprietary technologies to open standards. And I mean tomorrow, if you can. That involves addressing where and when data should be analyzed, security concerns, and the evolving relationship between the central IT function and OT roles within LOBs. Equally important, however, is to encourage your vendors, suppliers, and ecosystem partners to adopt open standards, participate in standards efforts, and do whatever you and they can do to stimulate and embrace standards-based technologies. That will do the most to save you money, headaches, and time when you begin to deploy IoT in production environments and provide you with a scalable foundation that will benefit you long term in your IoT journey.

This state of IoT creates something of a conundrum. On the one hand, we have a desperate need for technology convergence, simplification, and interoperability. We also need to rationalize disparate technologies around open standards and integrate them with legacy systems. On the other hand, IoT requirements and use-cases are diverse and still rapidly evolving. New devices, technologies, and methods are introduced

daily for which there may not yet be a standard. How would you resolve this conundrum?

Overabundance of Access Technologies

As noted above, we're all facing an overabundance of access technologies. This is further complicated by disparate devices and the tasks we want our various IoT solutions to accomplish. To even begin to decide among this variety of access technologies, you need to answer these basic questions:

- How many and which types of devices are in your network?
- Are these devices mobile or fixed, and how geographically dispersed are they?
- In which type of physical environment will these devices operate?
- How much data is being transmitted, and what bandwidth is required?
- How time-sensitive are the data transmissions?
- For battery-powered devices, what likely duration of operation is required, and how long should the battery last?
- What are the cost constraints?

Notice that these aren't technical questions, just basic business questions any LOB manager would want answered. Deciding on the right access technology is only the first step in designing IoT capabilities that will drive efficiency and yield actionable insights and better decisions. Before you have a working IoT solution, you'll also want to address the migration from proprietary technologies to open standards (yes, I know I sound like a broken record here), where and when data should be analyzed, security and risk assessments, and the evolving relationship between the central IT function and OT roles within LOBs.

Common IoT Framework

Are you overwhelmed yet? Fortunately, the industry quickly recognized that we can't go through each building block in a solution and ask basic questions like those above without a decoder ring. Thus, we've started to converge on a common IoT framework.[2] That IoT framework is not

just a slick marketing gimmick. It actually represents the way serious IoT players think about issues such as architectures, terminology, and logical blocks. Put another way, it's about using common designs, described in the same terminology, to refer to the same things. This isn't cast in concrete yet; it's only a starting point.

Such a framework can guide us on how to reduce the complexity of IoT technologies and solutions. It helps us determine which layers to abstract, where to focus on interoperability, and where to create open APIs as well as common and open standards. Since IoT is still evolving and will continue to evolve for years, we want a way to accommodate new innovations while ensuring that any new things can work with existing things. Otherwise, we'd all be in the position of reinventing everything anytime something changes. The IoT World Forum Reference Model is a good example of such an effort. (The IEEE IoT Architectural Framework is another example.) Its common framework drives interoperability across all IoT components: devices and controllers, networks, edge or fog computing, data storage, applications, and analytics. The model (Figure 10.2) organizes these components into layers and provides a graphical representation of IoT and all that it entails.

Equipped with such a reference model, the IoT industry has been focusing on three different standardization thrusts:

1. **Evolving existing horizontal standards.** As has been the case with many previous technology transitions, the robust standards of the IT world are now evolving to include requirements from OT and IoT. Dozens of interest groups in the IEEE, IETF, and other standards bodies are working on requirements for IoT, including time-sensitive networking (TSN) for cars or industrial control systems and safety; high-speed mobile communications among diverse things such as cars, trains, and other vehicles; or high-coverage low-speed networking technologies for low-power low-bandwidth sensors.
2. **Migrating specialized, proprietary, and semi-standard technologies to open standards.** As we've discussed, major industry players in manufacturing, transportation, and other verticals have historically implemented proprietary technologies or established standards around their own protocols and technologies. This often created conflicting standards, thus inhibiting interoperability and adoption. The IoT industry is working with the major industry standards bodies,

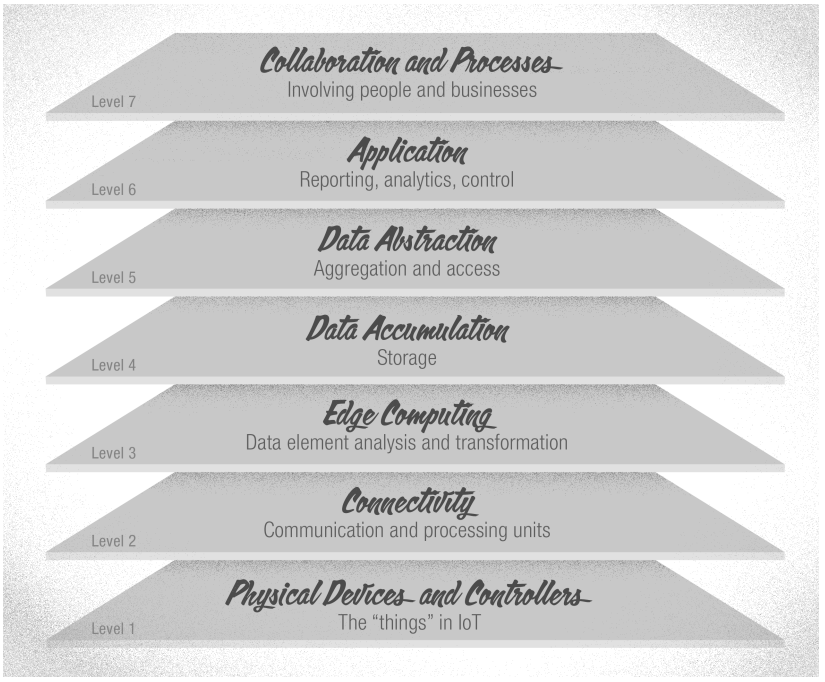

Figure 10.2 The IoT Technology Stack
Source: IoT World Forum Architecture Committee, 2015.

including ODVA, to avoid this issue by migrating to open standards while ensuring interoperability with legacy protocols.

3. **Creating consortia to address key pain points.** Major industry players are joining forces in new consortia, among them the Industrial Internet Consortium (IIC), the Open Connectivity Foundation (OCF), the OpenFog Consortium (OFC), and the OPC Foundation (for open platform communications).

IoT technologies are organized as a technology stack that moves up from physical devices at the bottom through data and applications, and finally processes. As I've mentioned, data analytics and vertical applications are key drivers for IoT. Most recently, I've seen an increased interest in data-in-motion and real-time/near-real-time data capabilities (think predictive analytics and fast payback scenarios), which are driving the latest interest in fog computing. One big challenge with such data

is that the data streams tend to age, which drives down their value very quickly; thus, the need to implement real-time analytics capabilities at the edge. Think about using data to identify and stop fraud as it's occurring. This isn't something you want to do hours, days, weeks, or months later. Besides the growing interest in fog-based analytics, the good news is that the industry is quickly adopting an open-source innovation model for both data storage and data governance, which should also speed data processing.

Finally, many of the challenges with IoT aren't technology-related but instead come from the industry's slow adoption and, often, resistance to change. One example of why creating common standards is so important: Wireless HART and ISA100, two different wireless standards focused on connecting sensors to the network. Both were derived from the IEEE 802.15.4 protocols, but each was created by a separate ecosystem of industry players and, as a result, is incompatible with the other. When my team came across these standards a few years ago, we thought we could help the industry converge on a common open standard for the wireless connectivity of sensors. That way, customers could easily choose among many sensor vendors and infrastructure vendors and not be locked into buying devices that only support a given standard. We went to both camps and proposed that we work on a plan to converge Wireless HART and ISA100 into one new open standard. Unfortunately, the idea was dead on arrival. I still hope that someday both standards will converge, but I don't see it happening in the near future. It will only happen when customers demand it and vote with their purchase orders.

Business-Relevant Standards Activities

I recently spoke with Max Mirgoli, executive vice president, World Wide Strategic Partnerships at IMEC, a world-leading research organization in nanoelectronics. He sums up the current standards situation this way: "With the advent of fast and simple connectivity, improvements in image sensing and other advanced sensing capabilities which can be tied together with simple yet powerful algorithms and apps, the IoT revolution has already began. We are starting to see early successes in smart manufacturing, autonomous connected cars, and smart grids, but the lack of convergence

on standards can slow the adoption. The good news is that pretty much all the major industry players recognize that without common standards, none of them will fully realize the economic potential of IoT. Thus, with the emergence of standards such as 5G, I am optimistic that the industry will band together to solve key technological and architectural challenges of IoT via common standards and interoperability."

I couldn't say it better myself. Standards efforts are important. With so much at stake in IoT, we want to avoid standards chaos and standards wars whenever possible. Remember the video industry's Betamax and VHS wars in the 1980s and 1990s? Or, even before that, the audiotape recording wars, when audio cassettes and 8-track tapes fought it out? Standards invariably benefit everybody. The same will be true with IoT, but even more so.

What follows is a brief summary of the main standards initiatives that are important to businesses embarking on IoT. This is by no means a comprehensive list. It's also subject to change as standards efforts emerge, depart, and evolve.

Horizontal Standards Efforts

- IEEE has kicked off a specific IoT initiative (see http://iot.ieee.org/). "IEEE has a long-standing track record of driving technology transitions through standards and interoperability. IEEE IoT Initiative is a multifaceted undertaking that brings together industry, academia, entrepreneurs, and investors," said Oleg Logvinov, who leads the industry engagement track for the initiative. He went on to tell me that "from the creation of the standard for an IoT Architectural Framework (IEEE P2413) to closing the gap between policy and technology development (IEEE Internet Initiative), IEEE is taking a very comprehensive and ambitious approach to fostering the creation of an IoT ecosystem based on open standards."
- International Telecommunication Union (ITU) Study Group 20 is developing IoT standardization requirements that will initially focus on smart city applications (see http://www.itu.int/en/ITU-T/studygroups/2013-2016/20/Pages/default.aspx).
- oneM2M Consortium (http://www.onem2m.org) is defining standards for a common M2M service layer to connect devices with M2M application servers. It targets business domains such as connected transportation, health care, utilities, and industrial automation.
- In both the AVnu Alliance and the IEEE, the industry is developing a set of standards around Time Sensitive Networking (TSN). "Time Sensitive Networking aims at building a foundation for more open,

easily accessible, and highly secure real-time control systems for the IoT," explained Georg Kopetz, member of the Executive Board at TTTech, an early pioneer of TSN. "For customers with mission critical applications, TSN offers real-time guaranteed latency, low-jitter and zero congestion loss for time critical traffic in converged networks," he added. As I discussed before, real-time analytics and apps are one of key drivers of IoT. That's why guaranteed network latency or delay that TSN offers is so important. TSN is enabling a standard-based approach to many use-cases from connected vehicles to the motion control applications on the factory floor.

Industry Consortia

- IIC (http://www.iiconsortium.org/) is working to accelerate IoT development and adoption in the industrial sectors to interconnect machines, business flows, intelligent analytics, and people at work. It has created reference architectures, established a range of innovation test beds, and is now identifying core standards, as well as gaps and requirements for future work. "The IIC has become the global consortium for Industrial IoT collaboration. With the membership of over 250 companies and 20 testbeds, the Consortium is evolving its Industrial Internet Reference Architecture and forging close collaboration with the Industrie 4.0 consortium," Paul Didier, Cisco's representative to IIC told me.
- OCF (https://openconnectivity.org/) is defining connectivity and interoperability requirements for connecting billions of devices. It is driving interoperability for device-to-device, device-to-infrastructure, and device-to-cloud communication by defining specifications, as well as creating open-source code and a certification program. This is a must-do to integrate billions of devices, sensors, and the data they generate into IoT solutions in a scalable way.
- OFC (https://www.openfogconsortium.org/), mentioned earlier, is developing an open-fog computing architecture for distributing computing services and resources close to users and endpoints to meet the growing demands for local computing in IoT. It will be releasing its reference architecture as this book reaches bookshelves.
- OPC Foundation (https://opcfoundation.org/) is leading the efforts on data interoperability, manufacturing processes, and equipment in the automation domain via its Unified Architecture. Thanks to its track record as an industry neutral forum, it is attracting new participants and expanding its scope across the entire technology stack. I expect it will continue to strengthen its role as the place where the industry gets aligned.

Industry-Specific Standards Bodies

- ODVA has been working tirelessly since the 1990s to champion open standards in the automation world and to migrate existing industrial automation standards to IP and Ethernet while ensuring interoperability with legacy protocols.
- ISA is tackling a wide range of standards issues, certifications, education, and training for the automation industry.
- PI, the umbrella organization for PROFIBUS (Process Filed Bus) and PROFINET (Process Field Net), is driving both sets of technologies.

It may seem as if we're sometimes taking two steps forward and one step back in terms of standards, but in general I'm very optimistic. The proponents of open standards clearly have momentum. Just visit the Hannover Fair, the largest industrial trade show on earth, and you'll see all of the devices proudly displaying their new standard Ethernet or wireless interfaces. The next step is for the customers to actually turn off the I/O interfaces and proprietary/specialized networks on their smart devices and start using such standards-based connections. So it's only a question of when, rather than if, open standards will become the norm in IoT.

New Technology Arrivals

Even as I was writing this book, new technologies emerged—and they will continue to do so. It also quickly became apparent that I could never include all of them and still finish the book. Instead, I decided to highlight a few that I consider the most important and far enough along to write about. What comes after, you'll have to discover on your own. That shouldn't be too hard. Just stay involved with your industry association and/or check out industry and IoT conferences and trade shows once a year or so.

Fog Computing

You've read my prior references to fog computing. Specifically, fog computing creates a platform—comprised of what we call a fog node—that provides a layer of compute, storage, control and networking services, and event stream processing between end devices on the ground and in cloud computing datacenters. Fog isn't a separate standalone architecture;

instead, it extends and scales the existing cloud architecture all the way to the edge of the network, as close to the source of the data as possible. The purpose is to enable real-time data processing and analytics of either large amounts of data or data in motion. The objective of fog computing isn't connecting devices differently. Rather, it's analyzing the data from the devices faster, with less latency and more efficiency. In effect, with fog computing we're putting data processing closer to the devices that generate or collect that data (Figure 10.3), and then analyzing it right there in real time.

A few years ago, Flavio Bonomi—founder and CEO of Nebbiolo Technologies, which focuses on the application of IoT technologies in industrial automation—led the definition (and naming) of fog computing with his team. When I asked him about fog, he summarized it well: "As we started to work on projects such connected vehicles, smart grids, and smart cities, we identified a common set of requirements for compact, scalable, well-managed, secured, and integrated networking, computing,

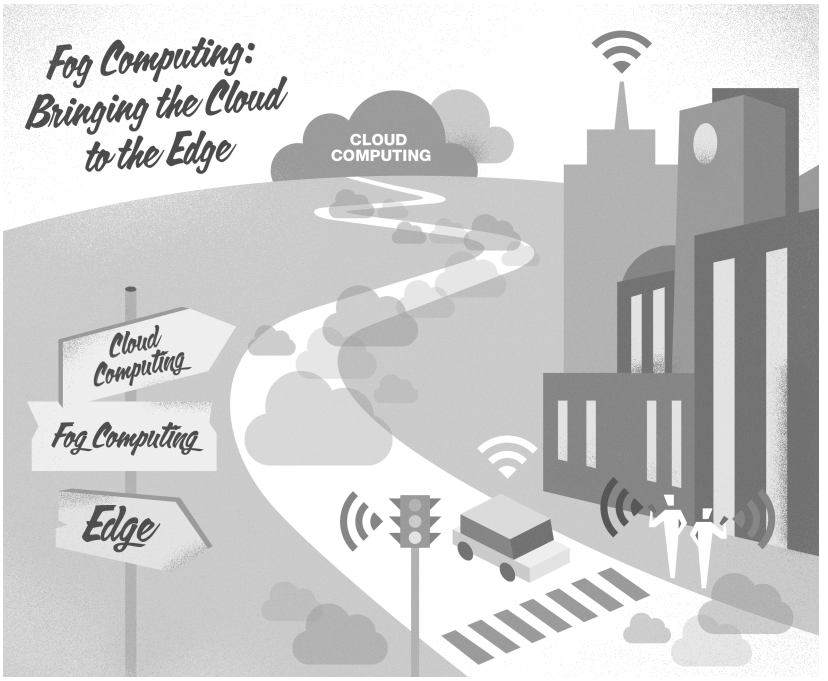

Figure 10.3 Fog Computing: Bring the Cloud to the Edge

and storage resources between the endpoints on the 'ground' and in more distant clouds. The term 'fog computing' was, in fact, naturally motivated by this need to bring more cloud-like capabilities 'closer to the ground.' In time, it became clear that fog computing actually facilitated the convergence of OT and IT and enabled new IoT use cases that required real-time capabilities, deterministic performance, physical security, and safety. Since it inherits elements from both IT and OT, fog computing naturally mediates between both domains at the various levels of the stack, from networking to security to the data level to the application level."

So what's the big deal about fog computing? At first glance, it doesn't appear to be all that different. In truth, however, it amounts to a distinct innovation. Fog computing (Figure 10.4) brings analytics and processing to the data. That's the difference, and it's a big difference. In the past, we always brought the data to where the processing occurred. That generally meant sending information to some distant central datacenter, which

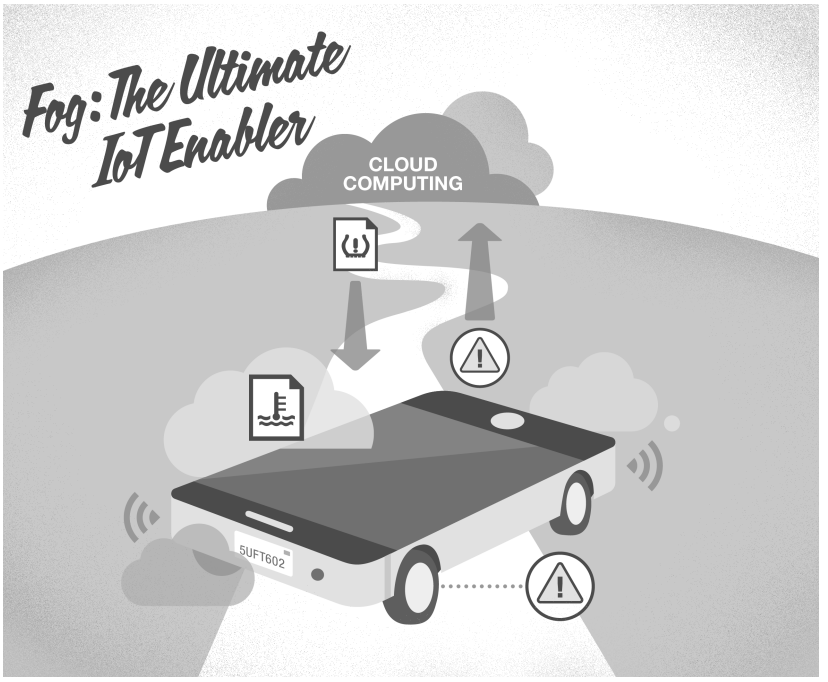

Figure 10.4 Fog Computing: The Ultimate IoT Enabler

added cost and significant delays. Now, with fog computing, we can scale the cloud and make it viable for real-time use-cases—the cloud and the edge can work together as an integrated system. Cloud software can send a policy to the fog node, requesting only certain types of data or only the exceptions to, say, a temperature threshold. The data is processed in the fog node based on this policy and only these exceptions, and the specific data requested is sent back to the cloud. The rest of the data is either stored locally in the fog or discarded.

As a result, we can convert the raw data collected from connected devices into useful information that can be acted on immediately—often in real or near-real time. When fog computing removes the latency from an IoT transaction, things can happen that fast. From there, we can also convert that information into valuable business insights through new applications, including real-time analytics and predictive context.

In short, fog computing brings:

- Near-real-time or real-time processing and analytics capabilities to the edge of the cloud
- Processing and analytics closer to the data and where they are used
- Much faster and more efficient analytics via a policy-based edge-to-cloud-to-edge system

Consider that the first stage of the Internet focused mainly on batch processing, wasn't time sensitive, and didn't use machines that consumed a lot of bandwidth. Now consider that even a single automobile can generate a huge amount of data and requires serious bandwidth—especially because that data is more time-sensitive and, therefore, even more important. (As an example, ask yourself how long you have to react if your car starts to overheat.)

Enter fog computing, which solves some of today's most common challenges, including:

- High latency on the network
- End-point mobility
- Loss of connectivity
- High bandwidth costs
- Unpredictable bandwidth bottlenecks
- Broad geographic distribution of systems and clients

As we've discussed throughout this book, fog computing is a key enabler of IoT, and it's driving an array of new use-cases in every area of life and industry—from retail to healthcare to oil and gas exploration and production. Preventive vehicle maintenance is one example. The sensors in each new connected vehicle generate up to two petabytes of data each year. It would be impractical and prohibitively expensive to send all of this raw data over the mobile network to the cloud for real-time processing. Fog computing turns these vehicles into mobile datacenters that can sort and index the data in real time and send alerts when action is required—for example, checking an overheated engine or filling an underinflated tire.

The industry has recognized the transformational capability of fog computing to enable a new wave of use-cases that weren't possible with cloud-centric implementations—hence the November 2015 creation of the OpenFog Consortium. "We formed OFC to accelerate the adoption of fog computing to solve pressing bandwidth, latency, and communications challenges associated with IoT, artificial intelligence, robotics, and other advanced concepts in the digitized world," OFC Chairman Helder Antunes told me. "Our technical workgroups are creating an OpenFog architecture that enables end-user clients or near-user edge devices to carry out computation, communication, control, and storage. And we plan to accomplish these goals in a collaborative manner, where interoperability between technology vendors is also ensured."

Blockchain Opens New IoT Possibilities

Blockchain has emerged as a technology that allows a secure exchange of value between entities in a distributed fashion by maintaining a continuously growing list of data records that are protected from tampering and revision. The technology first appeared on most IT radar screens a few years ago in the form of Bitcoin, a virtual currency that relies on blockchain technology to ensure its security and integrity. Although Bitcoin's future is uncertain, blockchain is a different story.

As the currency's underlying technology, blockchain is attracting considerable attention for its ability to ensure the integrity of transactions over the network between any entities. For example, I spoke with an energy company looking at blockchain to manage the interactions between

solar panels and the power grid. Automobile companies are considering the technology to authenticate connected vehicles in the V2V environment. Among the many other uses of blockchain being considered are the ability to trace the sources of goods, increase food safety, create smart contracts, and perform audits. Blockchain, it turns out, is a natural complement to IoT security in a wide variety of use cases.

Blockchain IoT implementations are still at a proof-of-concept stage, but standards are already starting to emerge. The Linux Foundation set up the Hyperledger Project, a partnership with several dozen major technology and financial players, to hammer out an agreement on open-source blockchain standards. For now, blockchain is presented as a sort of distributed consensus system ledger or database where no one person or entity controls all of the data. In effect, blockchain creates and stores a permanent or immutable log record of every transaction. As an emerging open standard, compliant variations of blockchain could enable products or solutions to offer different levels of control and programmable business logic via smart contracts. We'll just have to watch and see what happens.

According to Martha Bennett, principal analyst at Forrester Research, blockchain could be a transformational technology that changes the game in banking, IoT, and beyond. "Long-term, blockchain has the potential to revolutionize distributed computing. Looking at it purely from a technical perspective, many of the projects currently underway are laying the foundation for new ways of approaching distributed computing inside and outside of banking—doing for the storage and application layer what the Internet did for the communications layer. It's early days yet, and it will take time for all of the security, privacy, and scale issues to be addressed,"[3] Bennett commented.

This much we know now: Blockchain, which produces and saves a distributed log of any type of transaction activity, enables people to put their trust in a "trustless" transaction environment. It essentially eliminates the need for a central trusted intermediary between buyers and sellers or, in the case of IoT, between communicating things. In fact, blockchain could potentially eliminate the need for any intermediaries in most transactions. For those who want open, trustworthy IoT communications without having to rely on intermediaries, blockchain, especially "private" blockchain, could provide the answer and enable the

type of distributed IoT exchanges people have barely begun to imagine could be possible.

Machine Learning Enhances Real-Time Analytics

Like blockchain, machine learning is another important technology for IoT. It delivers a critical technology behind real-time predictive analytics, one of the key IoT use-cases. Machine learning has been around for years, but the recent advances in deep learning, especially supervised learning, have made it more valuable to IoT. Basically, with supervised learning, you can train the analytics system to improve its predictive accuracy—the more data on device operation, failure, and maintenance you feed into it, the more accurate the predictive analytics system becomes. Furthermore, although unsupervised learning has not evolved at the same pace and still has many open issues, it, too, is proving to be an invaluable capability for IoT. Think about zero-day attacks, where the hacker is exploiting a vulnerability in the software that is at that time unknown to the software provider. In such a scenario, since no data is yet available to train a classifier, such as a neural network, advanced unsupervised learning is starting to be used to detect such attacks.

Self-learning networks (SLN) are a great example of disruptive power of machine learning in IoT. In short, SLN is an architected solution combining powerful analytics with a wide set of machine-learning technologies (including cognitive learning from machine to machine) that enables networks to become intelligent, adaptive, proactive, and predictive. SLN has been architected with high scalability in mind: To that end, a wide set of machine learning algorithms are used at the edge of the network, which constantly learns network traffic patterns in order to build mathematical models.

Such models can then be used for a variety of purposes:

1. Prediction of application performances: by predicting the level of quality of service that IoT applications will receive from the network, it becomes possible for the network to anticipate and adapt accordingly.
2. As we discussed, security is known as one of the main challenges of our industry, with constantly evolving attacks that are becoming more and more pervasive and sophisticated. SLN makes use of

> machine learning to compute highly sophisticated models capturing normal baselines. Such models allow for the detection of advanced attacks, such as data ex-filtration, and denial of service attacks against the IoT network.

The SLN gets smarter as more events occur: Each node in the network performs modeling using machine learning, and learns constantly. Hosted on network edge devices and connected via advanced networking, SLN enables the network to both detect and respond much faster to problems.

"The concept of self-learning networks was born in 2012 while we were working on highly challenging problems for the IoT. Over the past few years, we faced a number of fascinating technical challenges, which led us to develop a highly novel and disruptive architecture and technology. We just announced the first product of a family of SLN portfolio called Stealthwatch Learning Networks for the detection of advanced threats. Without a doubt, many more SLN innovations applied to the IoT will emerge over the next few years that will considerably impact the IoT architecture, enabling a wide range of new services and capabilities," commented JP Vasseur, Cisco fellow and inventor of SLN.

Fog computing, blockchain, and machine learning are just three of the myriad technological and architectural shifts emerging around IoT. Stay tuned; many more are incubating, driven by the new challenges and the new opportunities IoT creates.

With its combination of open standards, interoperability, and new technologies, IoT is gaining powerful new capabilities and business models that will define winners and losers across industries. Already, savvy LOB managers are asking for open, IP-based IoT architectures. And companies like Cisco and Rockwell Automation are working with an increasing number of partners that have made the strategic decision to embrace open standards and evolve to the open IoT model. The paybacks are real; I saw them firsthand at last year's IoT World Forum in Dubai, where IoT early adopters presented their results.

The next chapter, "IoT State of the Union," definitely is not a recap of everything you have read so far. I bring in some new ideas and provide a glimpse into IoT's future, although I am not a futurists in any way, shape, or form.