CONTROL

TECHNOLOGY REPORT

# The Technologies Reshaping Industrial Network Infrastructure

Plan ahead for the accelerating convergence of IT and OT architectures

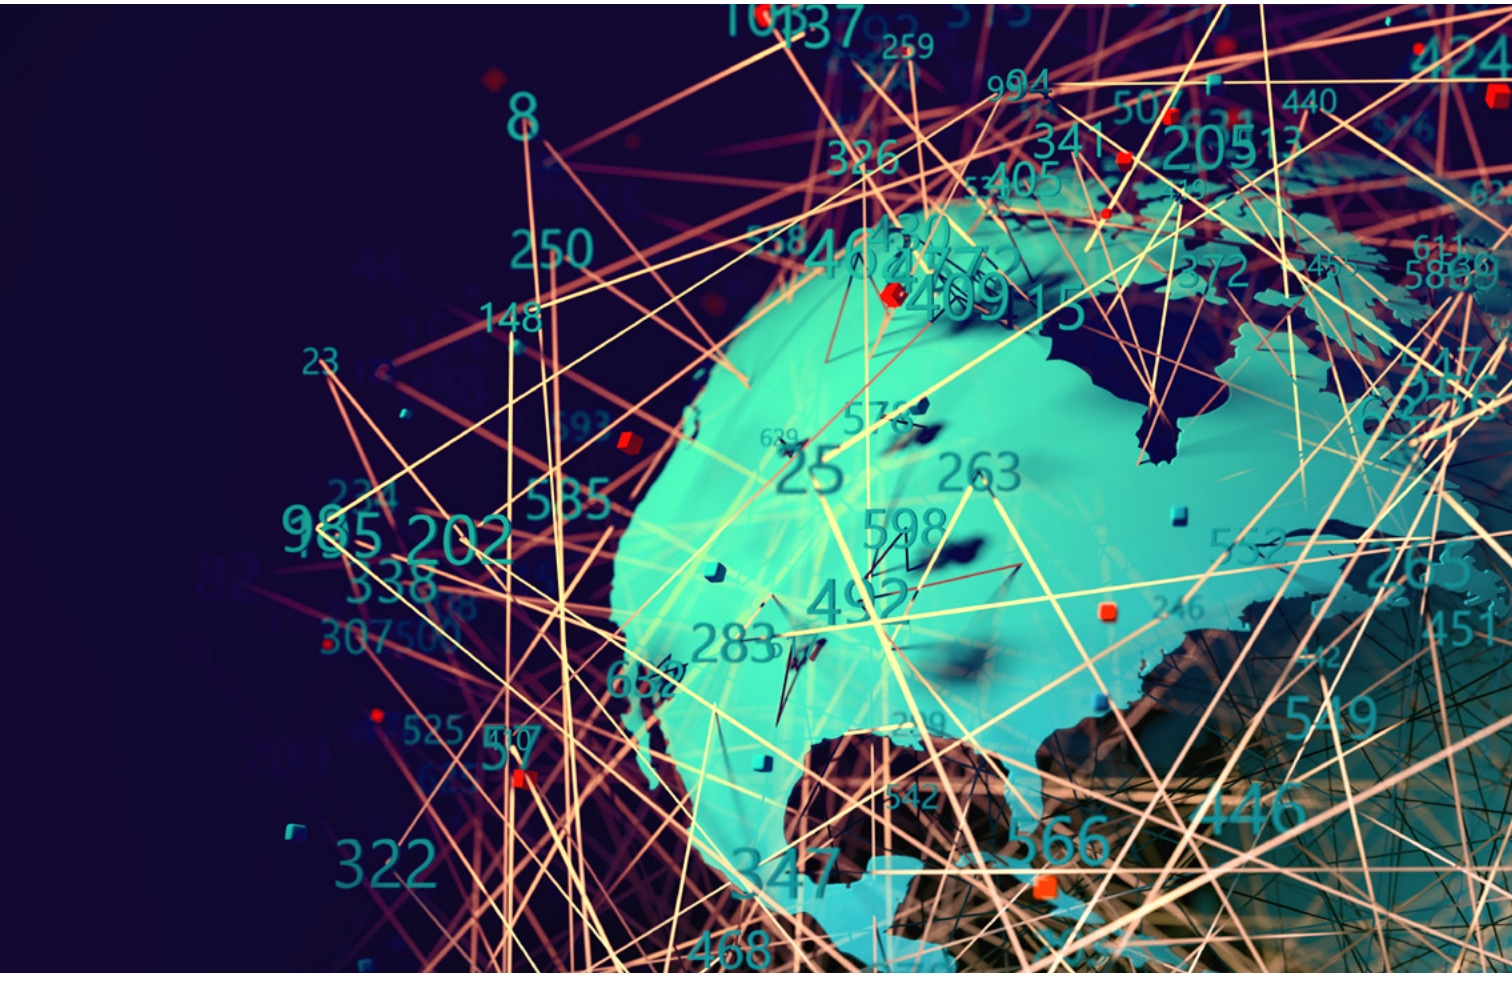

Sponsored by

PANDUIT®

## Is Your Infrastructure Friend or Foe?

As technology advances and Industry 4.0 becomes a reality, the demands placed on your network and electrical infrastructure will only increase. To stay competitive, you need plant floor solutions that improve production efficiencies, turn operations data into actionable insights, and scale for future growth. From cable management for control panels to fully integrated industrial networks, Panduit solutions turn connectivity into a competitive advantage. To read our white paper on Edge IoT deployments, visit **www.panduit.com/edgeiotdeployments.**

INDUSTRIAL NETWORK ZONING | CONNECTIVITY | DISTRIBUTION FRAMES | CABLE MANAGEMENT | WIRING DUCT

PANDUIT®
infrastructure for a connected world

**CONTENTS**



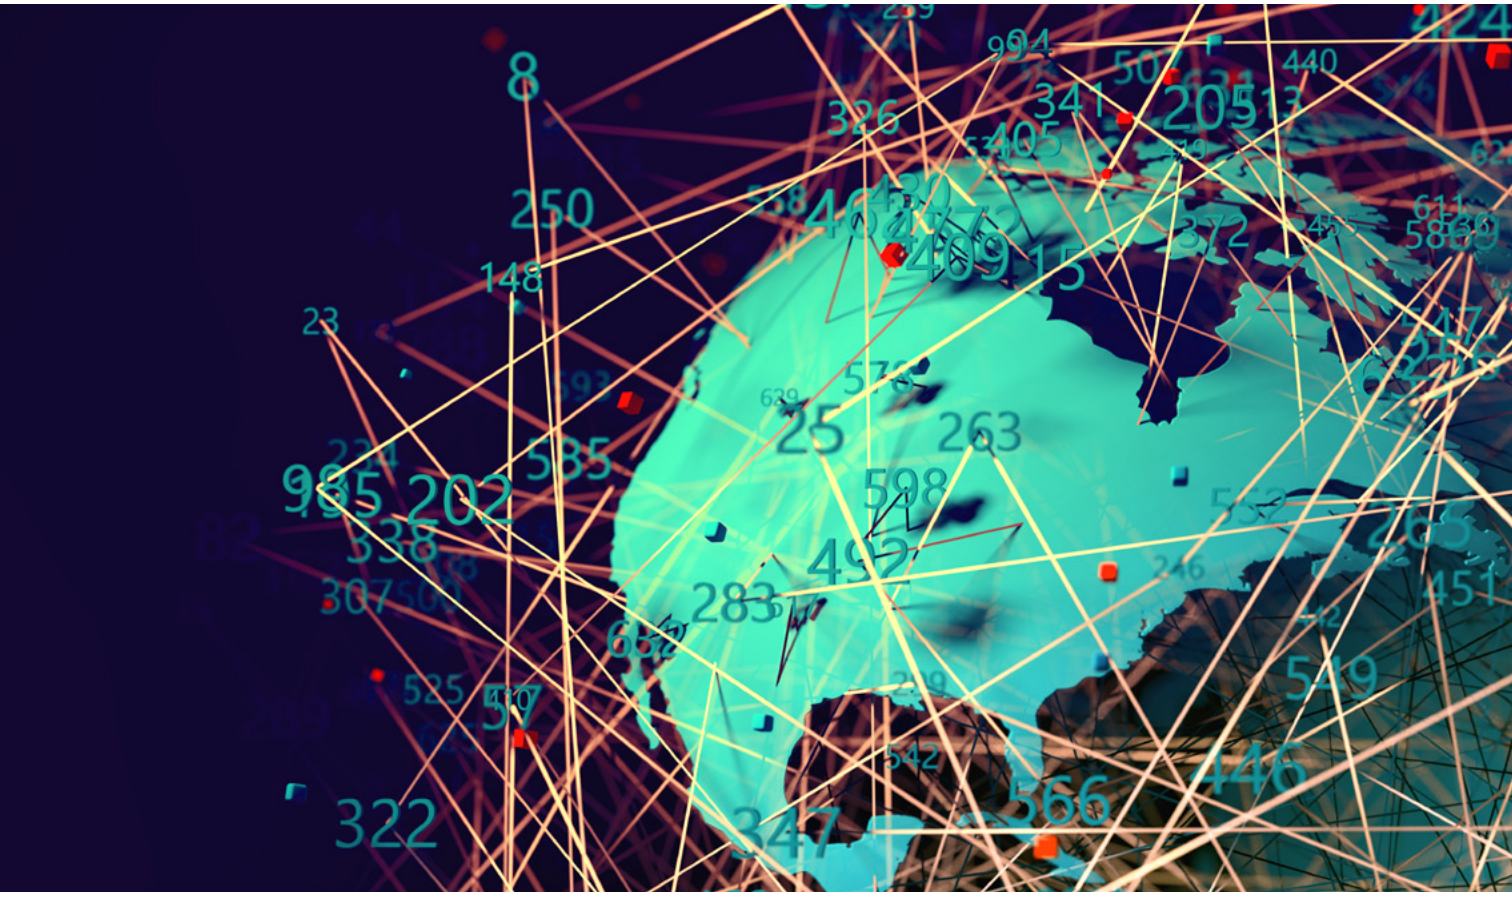

# Users sound off on industrial networking needs

Feeling ill-prepared for the demands of Industry 4.0, plants and factories are on the look-out for infrastructure technologies that can give business performance a boost

**By Keith Larson, Editor-in-Chief, *Control***

Out of sight, out of mind.

It's easy to imagine that much like a typical homeowner's benign neglect of the wiring, plumbing and HVAC systems that make modern life possible, industry's engineering and operations teams don't pay much attention to the industrial networking infrastructure that powers modern manufacturing. That is, until something goes wrong.

And while it's likely true that production problems due to network performance issues are quickly escalated for trouble-shooting, the majority of industry users say they are proactively on the look-out for new networking infrastructure technologies that can boost business performance. Such are the conclusions to be drawn from a just-completed survey of *Control* and *Smart Industry* readers on their attitudes toward—and adoption of—industrial networking infrastructure technologies.

**The majority of industry users say they are on the look-out for new networking infrastructure technologies that can boost business performance.**

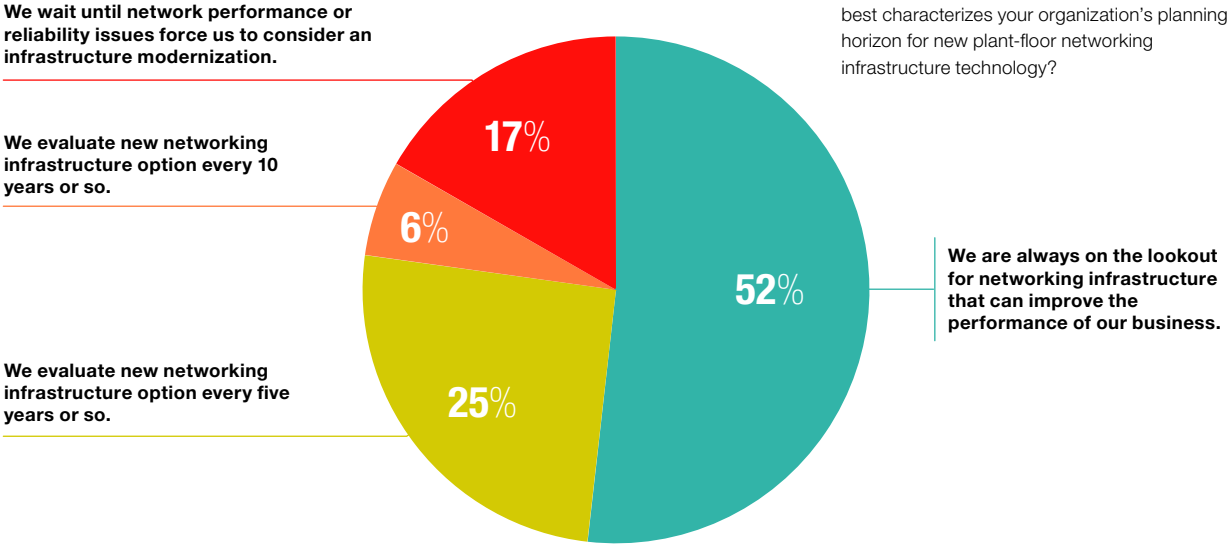


**Figure 1.** Which of the following statements best characterizes your organization's planning horizon for new plant-floor networking infrastructure technology?

## SURVEY METHODOLOGY & RESPONDENT DEMOGRAPHICS

In January and February of 2019, an email survey was sent to members of the *Control* and *Smart Industry* communities. A total of 114 survey responses were received, with industry, company size and job function demographics as follows.

### BY INDUSTRY

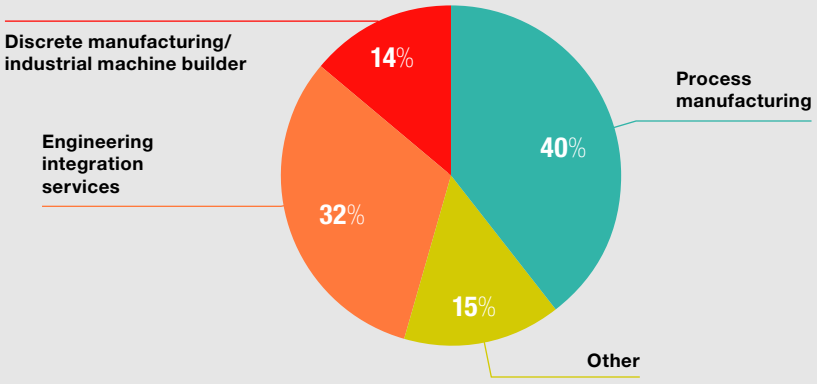


### BY NUMBER OF EMPLOYEES

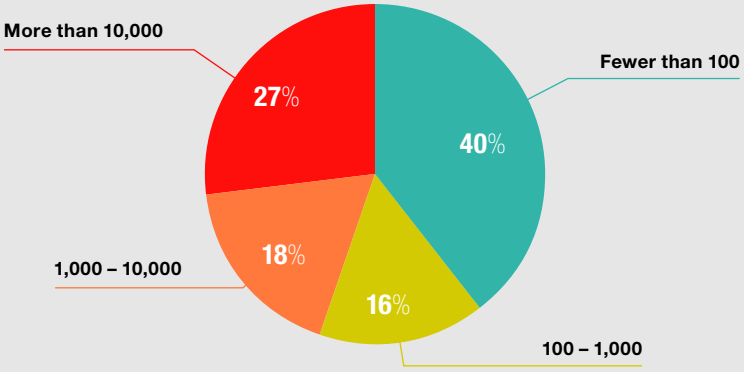


### BY JOB FUNCTION

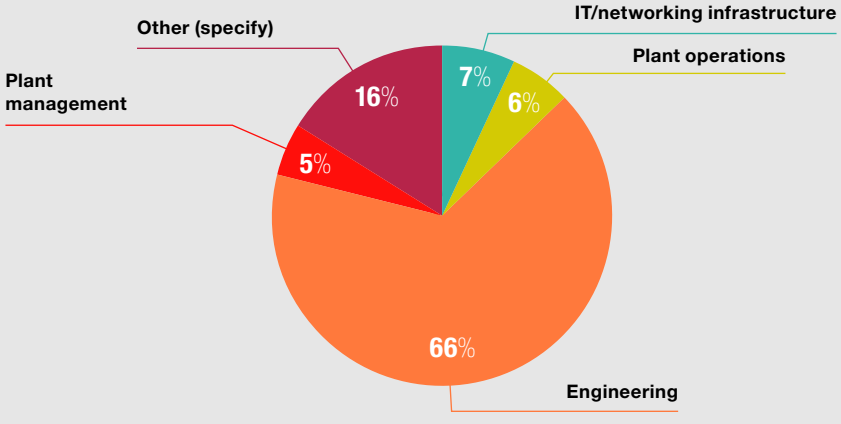

| | CRITICALLY IMPORTANT | SOMEWHAT IMPORTANT | NOT IMPORTANT |
|---|---|---|---|
| **Reliability** | 92% | 8% | — |
| **Performance** | 82% | 18% | — |
| **Conformance with open standards** | 51% | 42% | 7% |
| **Ongoing support services** | 42% | 50% | 8% |
| **Cost** | 29% | 67% | 4% |
| **Turnkey implementation** | 19% | 60% | 21% |

**Figure 2.** How important are the following considerations when selecting a network infrastructure solution for manufacturing/production areas?

As indicated in Figure 1, a full 52% of survey respondents indicate they are "always on the lookout" for networking infrastructure technologies that can improve the performance of their business. Another 31% of survey respondents say that they periodically evaluate their network infrastructure options every 5 or 10 years (25% and 6%, respectively). On the other end of the spectrum, 17% of respondents are strictly reactive, waiting for network performance or reliability issues to force them to consider an infrastructure modernization.

Meanwhile, those surveyed universally agreed that reliability and performance are far and away the most important considerations when selecting a network infrastructure solution for use in manufacturing/production areas (Figure 2). Conformance with open standards and ongoing support services were the next most important criteria, followed by cost and turnkey implementation.

We also asked a series of questions specifically about wireless networking infrastructure in production environments, starting with current technology usage (Figure 3) and the use cases driving the deployment of wireless. Not surprisingly, Wi-Fi was the most

## Reliability and performance are far and away the most important considerations when selecting a network infrastructure solution

| | YES | NO | UNSURE |
|---|---|---|---|
| **1. Wi-Fi (wireless LAN)** | 84% | 13% | 4% |
| **2. Commercial mobile (4G/LT E)** | 35% | 54% | 11% |
| **3. WirelessHART (low-power mesh)** | 29% | 60% | 10% |
| **4. Private mobile (4G/LT E)** | 21% | 67% | 12% |
| **5. DAS (distributed antenna systems)** | 14% | 75% | 12% |

**Figure 3.** If wireless networking is used in your facilities' production areas, which of the following technologies are in use?

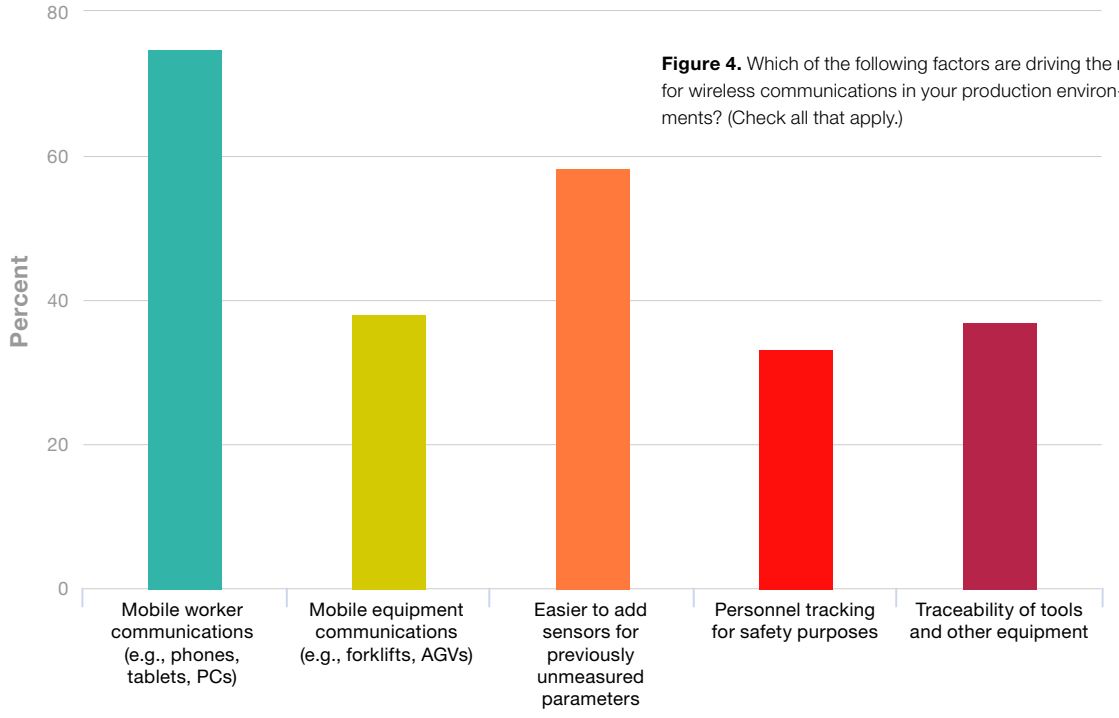


**Figure 4.** Which of the following factors are driving the need for wireless communications in your production environments? (Check all that apply.)

commonly used (84%) followed by commercial 4G/LTE (37%) and WirelessHART (29%). Private 4G/LTE and DAS, or distributed antennae systems, were reported in use by 21% and 14% of respondents, respectively.

In terms of the use cases driving the deployment of wireless infrastructure, mobile worker connectivity and the ability to easily add sensors for previously unmeasured parameters were most often cited (Figure 4). Mobile equipment communications (as for automatic guided vehicles, or AGVS, and forklifts), traceability of tools and other equipment, and personnel tracking for safety purposes all made the list of important use cases for more than 30% of respondents.

Interestingly, when it comes to decision-making on the use of wireless in production areas, it's most often up to engineering leadership, but IT and operations also carry sway in some organizations (Figure 5).

The final pair of questions in our survey focused on the evaluation and pursuit of key *emerging* wireless networking technologies, including 5G and the recently completed Wi-Fi 6 specification. While both of these options are under evaluation by about a third of survey respondents' organizations, implementation of Wi-Fi 6 is being actively pursued by nearly twice as

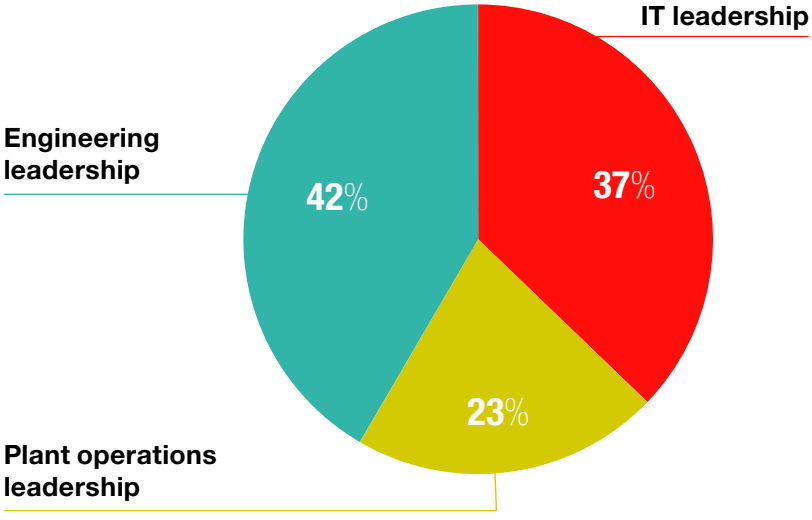


**Figure 5.** Within your organization, who is the key decision-maker on the use of wireless networking technology in production environments?

| | NOT CONSIDERING | UNDER EVALUATION | PURSUING IMPLEMENTATION | UNSURE |
|---|---|---|---|---|
| **Wi-Fi 6** | 31% | 29% | 7% | 33% |
| **5G commercial** | 38% | 29% | 4% | 29% |
| **5G private** | 42% | 23% | 1% | 33% |
| **Citizens Broadband Radio Service (CBRS)** | 52% | 11% | 3% | 34% |

**Figure 6.** To what extent has your organization pursued the following emerging network technologies?

many—albeit still relatively few—respondents' organizations as 5G (Figure 6).

Finally, we asked our survey respondents whether they felt their facilities' current wireless networking infrastructure was up to the task of handling the growing demands of Industry 4.0. A slight majority of respondents ranked themselves a middling "somewhat confident" (Figure 7).

Meanwhile, 1 in 5 were either "very confident" or "not confident at all." Two percent reported their organizations have no intentions of using wireless, while 6%—presumably with tongue firmly in cheek—checked the box for "What's Industry 4.0?" ☐

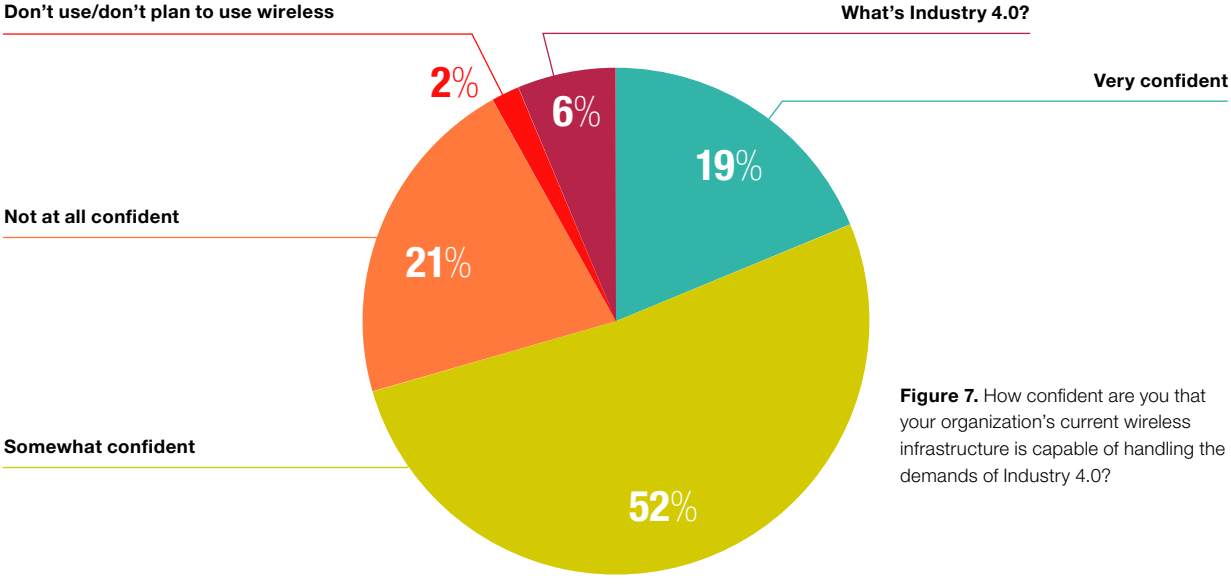


**Figure 7.** How confident are you that your organization's current wireless infrastructure is capable of handling the demands of Industry 4.0?

# Plan two steps ahead to support network infrastructure needs

*As industry builds towards a digitally transformed version of itself, future success will often rely on the foundation of physical network infrastructure put in place today. That means new wireless connectivity to enable mobile workers and coordinate the movements of autonomous intelligent vehicles (AIVs) must work together with the physical network infrastructure that has long been the backbone of industrial connectivity.*

*Because so much is riding on that physical network infrastructure, Lindsey Parker believes it's important that individuals charged with designing and implementing solutions for this realm understand the long term implications of their decisions. Physical cabling is often expected to last a couple decades or more, says Parker, manager of industrial network business development for Panduit, so it's critical to satisfy today's networking infrastructure needs with an eye to how those needs and requirements will change down the road.*

Smart Industry *recently caught up with Parker to learn what factors are most important to consider in an industrial network infrastructure solution, and what steps can be taken to future-proof those investments for the bandwidth-hungry applications that are sure to come.*

Lindsey Parker is industrial network business development manager for Panduit.

**Q:** Panduit has long been a leader in providing the physical infrastructure for digital networks as well as electrical power for plant-floor and production environments. What have you found are the most important characteristics that industrial end users and OEMs seek in a network infrastructure solution?

**A:** First off, don't sacrifice performance and reliability over price and availability. Second, choose partners that understand the unique requirements of industrial environments. Third, plan not just for what you need today but what you'll need tomorrow and the day after as well.

It's important to know, for example, what type of media is suitable for your environment. You don't always need an IP67 rated enclosure or an M12 connector, but in our world, sometimes you do. Panduit references the TIA 1005-A standard for industrial environments, which provides guidance through mechanical, ingress, chemical/climactic and electromagnetic (MICE) considerations. Knowing when to invest in hardened cabling infrastructure and when to use commercial grade will save you money in the long run while ensuring your infrastructure can withstand its surroundings.

Also, remember that while the network infrastructure will only represent 7-10% of the spending on a given project, it's got to be properly configured. That means the choosing the right topology to provide appropriate performance and resiliency. 'Should be good enough' just won't cut it. Rather, consider following a proven approach such as the Converged Plantwide Ethernet (CPwE) reference architecture that Panduit worked with Rockwell Automation and Cisco to define

and validate. Using that architecture as a guide, you can know exactly what performance to expect from your completed solution.

**Q:** While physical cabling certainly remains the backbone of plant-floor automation and information networks, how has the infrastructure equation changed with the increased performance of wireless networks such as those defined by the WiFi 6 standard and 5G from the telecommunications industry?

**A:** High-performance wireless is an important technology that can deliver new capabilities to industry. But it's important to realize that it doesn't replace physical network infrastructure. Rather, the latest high-speed wireless standards such as WiFi 6 actually require a lot more access points in order to deliver that higher speed and bandwidth. And more access points mean more cabling!

That same equation applies to the 5G wireless infrastructure being promoted by the telecommunications industry. But with the WiFi 6 vs. 5G decision, the biggest question is not performance but whether you want to own your network outright, or enter into a contract to have that infrastructure managed on your behalf.

Another implication of the latest wireless standards is that they all but require Category 6A cabling infrastructure, which supports 10-Gb Ethernet. Many plants operate today operate on Category 5e cabling, which is only 1-Gb Ethernet capable. Also, since physical infrastructure is likely to outlive the radios that are connected to it—consider running a second, redundant Cat 6A cable while you're at it. Think two steps ahead so that you're better prepared for what's to come.

**Q:** What types of applications are made possible by the performance of these new options?

**A:** The biggest impact of these new wireless networks is around workforce mobility—basic information access first, then

Wireless apps need physical infrastructure, too! When planning for future network infrastructure requirements, remember that the latest high performance wireless networks need extra access points as well.

layering in really cool stuff like augmented reality. With AR, you see major improvements in employee training and the ability to more effectively utilize a remote workforce. A network engineer in Chicago can work with someone unskilled in networks in a plant in Iowa to troubleshoot a problem. A production operator can learn on the job using virtual work instructions instead of taking a training module in front of a computer screen and then trying to remember what it said out on the production floor. People expect to work how they live, and wireless communication is in our DNA now. We need it for productivity, safety and communication.

Other technologies like robotics and AIVs will really benefit from wireless technology, too. 5G is expected to be the catalyst for autonomous vehicles on the roads, and I think you can expect to see that be the case in the plant, too. And obviously, the whole IoT world is founded in wireless communications. We'll see an abundance of new sensors that are lower cost and easier to deploy and configure.

**Q:** What advice can you offer to industrial networking decision-makers seeking to pursue these speed- and bandwidth-hungry use cases?

**A:** Don't jump right in. You really need to take a look at your existing physical infrastructure to understand what your network can support in its current state. If it resembles the cabling that's out there in most of the industrial world, you're probably not ready to plug in WiFi 6 access points tomorrow. Bring in the right people to help you plan out the physical layer that will be the foundation for all of this new technology and then make sure it gets installed properly. Once you've done that, your investments in high-performance technology will be able to shine. ◻

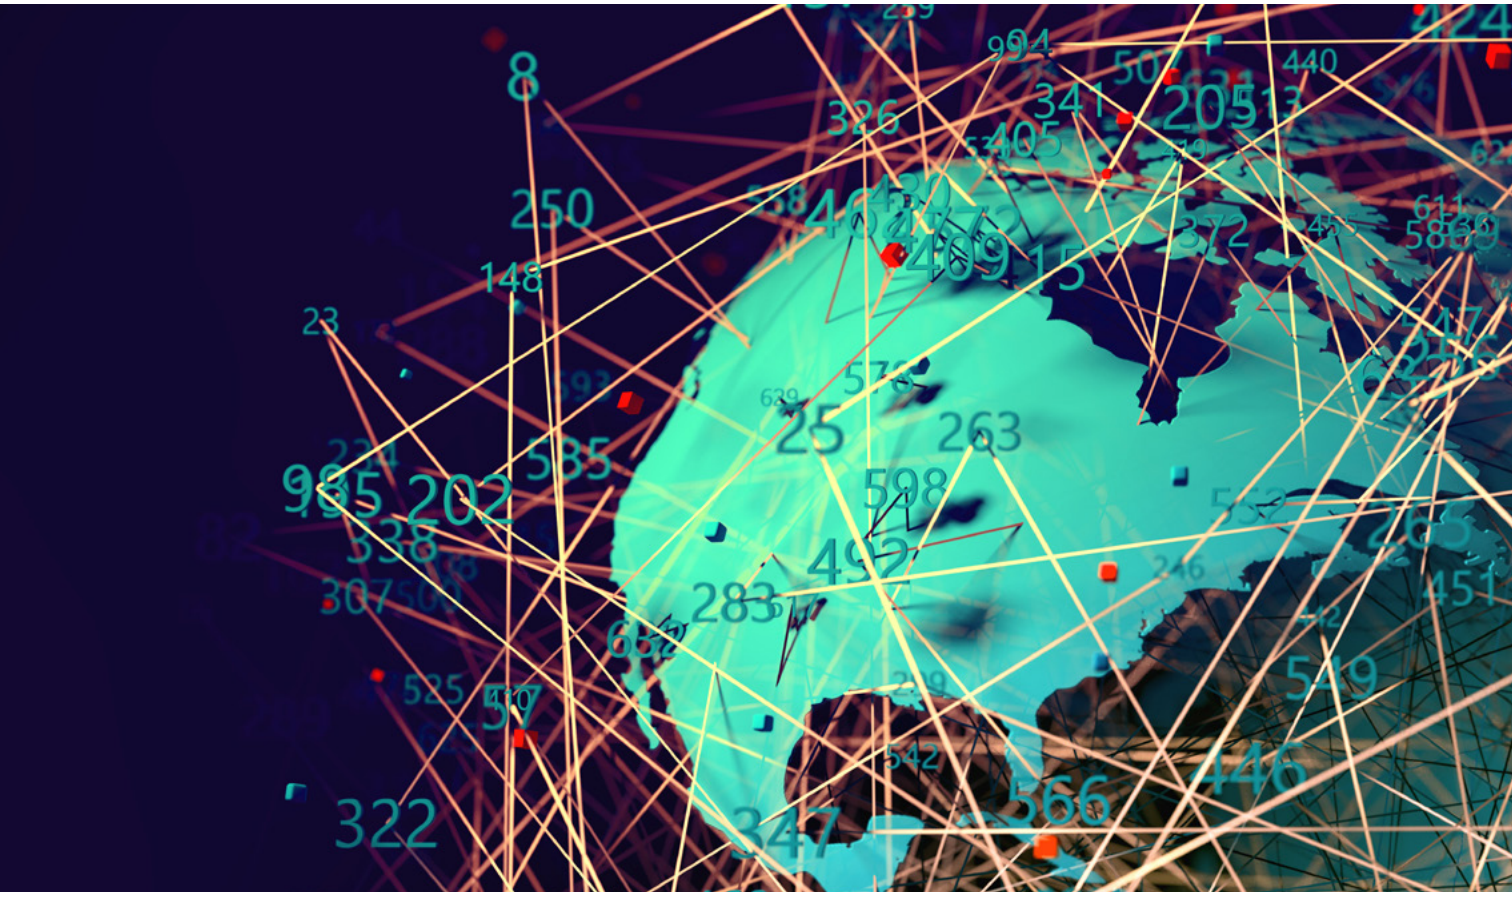

# IIoT, cloud computing changing control system architectures

Technologies are changing our view of the venerable Purdue model

**By Ian Verhappen, Senior Project Manager of Automation, CIMA+**

The Purdue Enterprise Reference Architecture incorporated in the ISA-95/IEC 62264 standard, on which the majority of control system architectures and subsequent standards including wireless, cybersecurity, safety, etc. are based, originated in 1989. Despite being in use for almost 30 years, many people still believe that it is based on physical layers, when it actually defines the functions to be performed at each level of the architecture. At the time the model was developed, and in most cases today, it is still true that form follows function, and the various pieces of hardware tend to correlate closely to their assigned function. The IEC 62443/ISA-99 cybersecurity zone and conduit concept also tends to encourage the maintenance and separation of each of the function-based layers.

With the changes in processing and computing capability we've seen at the different levels of the enterprise, particularly Level 1, and the introduction of cloud-based systems, it is my understanding that ISA-95, as part of their regular review of the document, is revisiting the architecture model, with particular emphasis on Level 1 and Level 0.

Another ISA standards body, ISA-112 SCADA Systems, also needed an architecture model on which to base their work. The present version of this model, which adds more granularity to the ISA-95 model, is shown in Figure 1.

**The virtualization of systems is changing system architectures once more, with the biggest impacts at the top and bottom.**

When creating this model, ISA-112 deliberately chose to use letters to show the different layers, in part to avoid confusion with the Purdue model (shown for reference on the side) but also to help the committee relate the physical equipment against the function(s) that equipment needs to perform.

In general, layers A through D will tend to be at the remote site, which could be anything from a single point and RTU to a remote compressor or pump station complete with its own "mini" control system with wireless SCADA connections to associated well pads, isolation valves or remote storage facilities, thus making "site n" a small SCADA system, or at least a data concentration site on its own.

Similarly, levels F and G identify the typical SCADA components that reside on the central SCADA server(s), typically in the main control center. Alarms and Historian have been identified as two typical databases residing at this level,

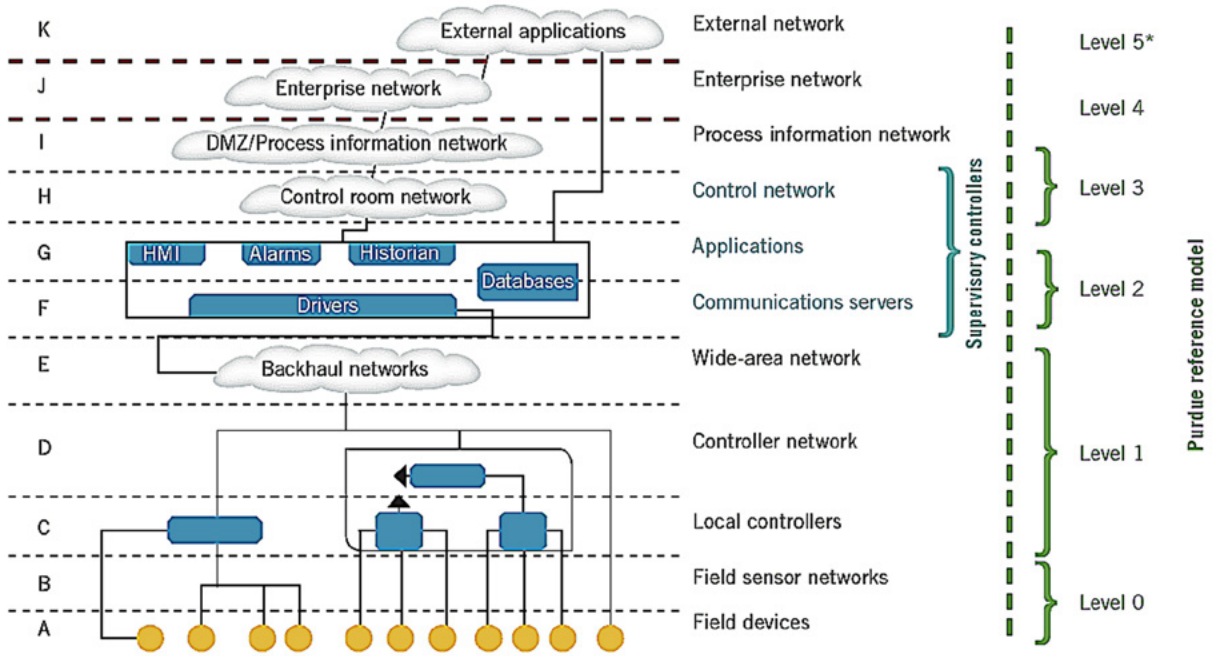


**Figure 1.** In this model by the ISA-112 SCADA Systems standards committee, letters are used to label layers to avoid potential conflict with ISA-95 and other similar models. Routers and firewalls between layers are not shown, nor are other system-specific servers, applications and workstation. Remote-hosted external applications (cloud) could not be configured to attach to devices at any level with appropriate firewalls, tunneling and routing.

** Note that although this diagram shows a Purdue level 5, the true Purdue model only has levels 0 to 4 because it did not anticipate external applications.*

though as indicated by the "database" box on the right, they are by no means the only ones; they are just the ones that the committee believes require particular attention since, from a SCADA perspective, they will have some unique constraints and items to be considered when developing a system.

The other addition to the proposed SCADA model is the concept of cloud computing, presently shown as the "external applications" cloud at the top. Though a link is only shown to the databases at the SCADA server, there is the potential to link to elements at any level, with, of course, the appropriate cybersecurity protocols.

Lastly, the red lines on either side of level J are intended to show the clear demarcation between the OT (SCADA related systems), IT and public or external networks as a reminder to pay particular attention to the cyber requirements when crossing between different layers and systems.

The virtualization of systems per Open Process Automation Forum (OPAF) and arguably IIoT, is changing control system architectures once more, with the biggest impact at the top (nonexistent Level 5 at the top of the model) and again at Level 1, with basic regulatory control moving closer to the process itself. Because more functionality in these models will reside in software versus the hardware-based representation, the case can be made that the function-based reference model will become even more important since the physical hardware could potentially be flattened into fewer layers residing in the cloud and a couple virtual machines for the hardware above the sensor layer(s). ◻

# The promise of single-pair Ethernet

New standards tackle the 'last mile' of connectivity for manufacturing automation and industrial control

**By Bob Voss, Senior Principal Engineer, JECIC Corporate R&D Center, Panduit**
**Jeff Beller, Business Development Manager, Industrial Network Infrastructure, Panduit**

Legacy manufacturing and industrial control networks must evolve to address the demands of Industry 4.0 and IIoT. Single-pair Ethernet promises to be the enabling technology that allows for cost-effective migration from many legacy protocols to one common protocol and addresses the need for a reliable, secure infrastructure providing high-bandwidth communication, power and control to edge devices. Single-pair Ethernet (SPE) extends the network to incorporate the "last mile" of connectivity creating a seamless Ethernet TCP/IP network fabric for the enterprise, from cloud to edge. SPE technology will help to build out the necessary foundation so enterprises can better achieve their smart manufacturing objectives and digitally transform their businesses. Its availability to the marketplace is projected to begin with early adoption in 2021.

## THE STATE OF THE NETWORK

More and more manufacturing and industrial equipment and devices are being connected to networks. Today's OT networks are a composite of Ethernet and legacy fieldbus protocols. What does the industrial network landscape look like? In terms of new installed nodes, HMS concluded that for 2018 industrial Ethernet had surpassed traditional fieldbuses for the first time, and this development continued in 2019. Industrial Ethernet continues with a steady growth rate of 20% and now makes up 59% of the global market, an increase of 7%. Globally, EtherNet/IP is the largest industrial Ethernet network with 15% of the market followed closely by PROFINET at 14%. Other significant Ethernet technologies include EtherCAT, POWERLINK and Modbus-TCP all of which are showing steady growth.

Most notably, 2019 was the first year of decline for new fieldbus nodes, declining by 5% compared to 6% growth in 2018. Fieldbuses in the aggregate now account for 35% of the global market, a decline of 7% compared to the 7% growth seen for industrial Ethernet.

Industrial networks evolve over time and this evolution has seen the gradual "cooling" of fieldbus deployments driven by the stronger growth of industrial Ethernet. As Andres Hansson, chief marketing officer at HMS has noted, "The transition to industrial Ethernet continues and is driven by the need for high performance and the need for integration between factory installations and IT-systems/IIoT applications." The arrival of SPE promises to further accelerate this transition.

**"The transition to Ethernet continues and is driven by the need for high performance and integration with IT/IIoT systems."**

## MEETING THE WORKFORCE CHALLENGE

The advent of SPE technology is very timely if seemingly unplanned. Manufacturers are being both pushed and pulled to transform their network infrastructures. Industries are at a crossroad where legacy systems are at, or are reaching, end-of-life while at the same time the workforce versed in supporting these systems are setting sail into retirement. Forward-thinking organizations have put plans in place to transfer and retain this knowledge. But for many it is unlikely that the collective tribal knowledge of this experienced workforce will be effectively transferred to the new and inexperienced one coming in.

As a stop-gap measure, some manufacturers have retained service organizations to support their operational technology (OT) networks. But this still leaves many manufacturers vulnerable and at risk. This is the push. The pull is the positive business outcomes that smart manufacturing and digital transformation delivers. Standing in the way of this, in part, are the machines and field devices connected with proprietary and legacy network protocols that that are not easily synchronized to 'uptake' into the enterprise systems for insights and action. SPE provides a protocol that the incoming workforce is already familiar with and presents a clear and cost-effective migration path.

## SPE OVERVIEW

For decades, Ethernet cabling has used four twisted pairs of copper conductors to carry data and power. But recent developments have changed all that. As a breakthrough technology and paradigm shift, Ethernet can now be deployed over a single balanced twisted-pair, thus single-pair Ethernet, or SPE.

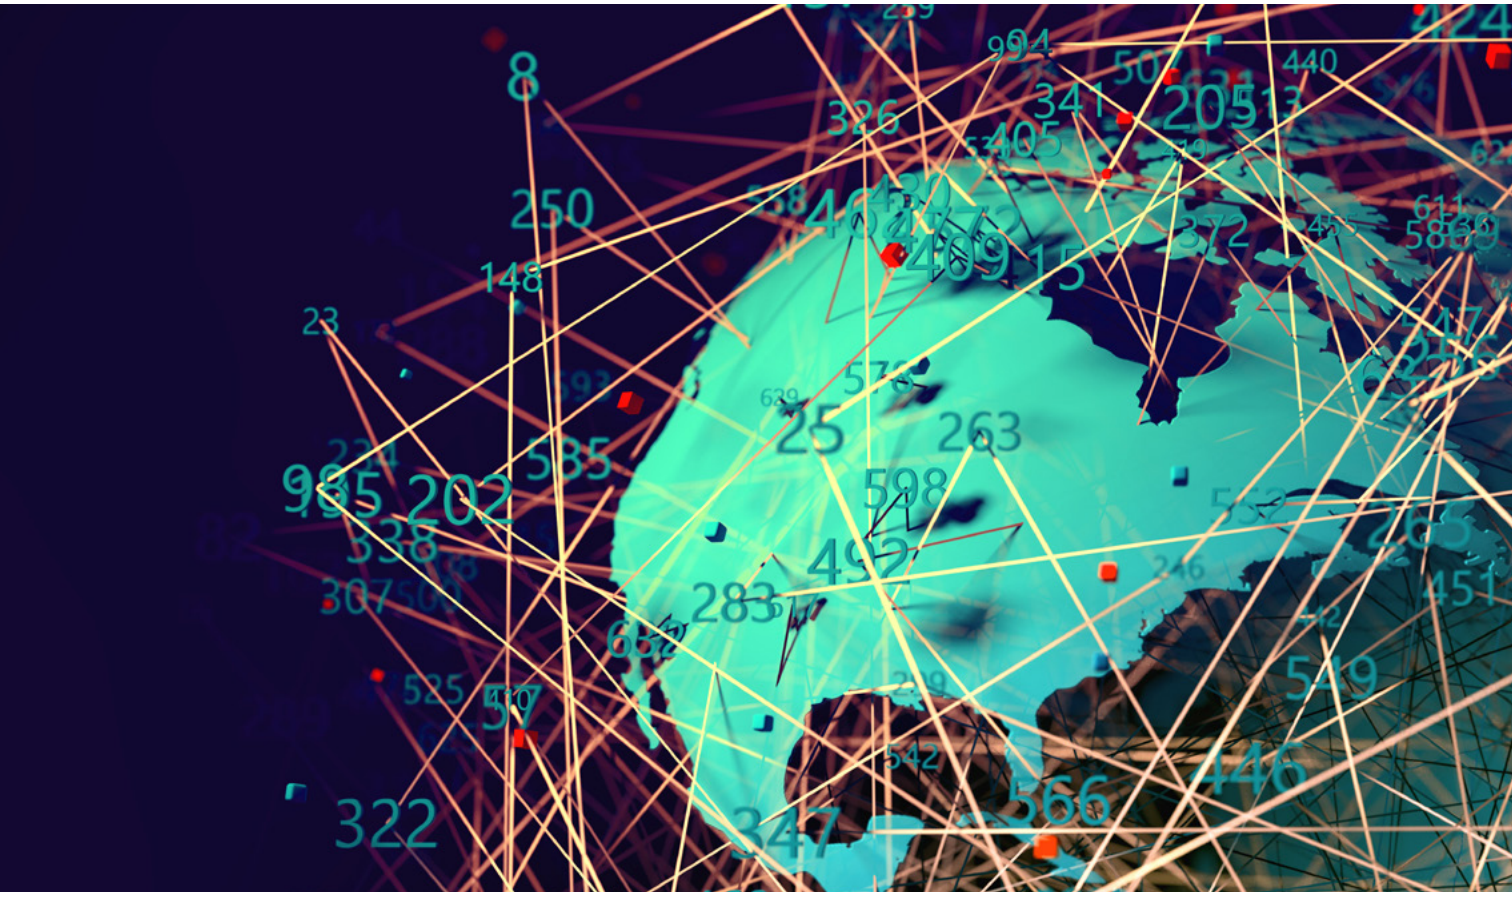

## Industry sees SPE technology as a means to achieve a single converged network—cloud to device—and displace the legacy systems that are prevalent.

In recent years, the automotive industry became the first to develop and adopt SPE, in this case for in-vehicle networks. This innovation was driven by the increasing number of applications and complexity of features being added to cars. With each feature introduced the number of subsystems proliferated along with much higher bandwidth requirements. Current bus topologies and a fragmented network architecture deployed in the vehicles were ill-equipped to meet these demands. The industry needed to converge the growing number of disparate systems and provide higher bandwidth—and all within very small spaces. Reducing the wire count was a driving force. The industry responded in 2015 with the IEEE 802.3bw standard for 100BASE-T1 for in-vehicle networks. Moreover, the supplier ecosystem worked with the automakers to develop and commercialize the technology for the benefit of the industry.

The 100BASE-T1 standard is the first SPE standard. It provides for 100Mb/s transmission over a single unshielded twisted-pair cable up to 15 meters in length. By superposition, the physical layer, or PHY, is full duplex so both the data send and receive transmission is over the same pair of wires; this is the engineering breakthrough that other application domains will leverage as well. Consider that standard Ethernet has a dedicated pair of wires for each direction.

In order to meet the demands for more bandwidth, industry further evolved and enhanced the technology to provide for 1 Gb/s operation over a single twisted-pair copper cable in an automotive application (IEEE 802.3bp-2016). In addition to converging and achieving high-performance in-vehicle networking, the technology enabled the industry to reduce component costs and reduce the overall weight of the vehicle, improving fuel economy.

The automotive industry's success with SPE was the impetus for manufacturing, process and control system industries to investigate the technology. Similarly, these industries saw the technology as a means to achieve a single converged network—cloud to device—and displace the legacy systems that are prevalent. Here, higher data bandwidth and power delivery requirements are inexorably linked. Pursuant to these goals, the IEEE 802.3cg Ethernet Task Force group was formed to study fieldbus technologies and create a standard for SPE to address the needs of OT networks for manufacturers and process industry. SPE is proposed to address the needs of Level 0 of the Purdue industrial control system hierarchy model. It is here that Ethernet is generally not present and where instead the proprietary and legacy fieldbus systems provide the connectivity from the control system (Basic Control Level 1) to the machines and field instrumentation.

The challenge and opportunity to transform OT networks is enormous. The fieldbus landscape is very fragmented, and variants include AS-Interface, BACnet, CAN, CANopen, CC-Link, ControlNet, DALI, DeviceNet,

Foundation Fieldbus H1, HART, Interbus, IO-Link, LonWorks, Modbus-RTU, Profibus D, Profibus PA, and others. There is no overarching standard resulting in far too many fieldbus variants. Oftentimes, facilities will have more than one fieldbus in operation to support. Despite the recent contraction of new fieldbus nodes as observed by HMS, there remains countless devices connected via fieldbus topologies. Consider that mid-range estimates suggest that globally there are 80 to 100 million new device connections per year. In addition, there remains a plethora of 4-20mA instruments hardwired to analog I/O.

The fieldbus protocols are serial communications usually operating over shielded and twisted two-wire cable. In general, the fieldbus systems provide very modest data rates and reach. Across the universe of fieldbus, the bandwidth ranges from 1200 baud to 12Mb/s. The longest reach ranges from 400m to 5,000m. For many of the systems, the data rate falls increasingly to very low levels as the maximum reach distance is approached. In determining the requirements for Ethernet, the new solution needs to cover a range of rate and reach available with fieldbuses with a single design**.** The 802.3cg working group concluded that the new SPE standard would be for 10Mb/s data rate and 1,000m (1km) reach to address nearly all fieldbus applications and enable a standard media access control, or MAC.

In general, no device power is delivered by the fieldbus connection. Therefore, local DC power supplies are required near the device to meet its power requirements. Behind the DC power supply are many AC components to convert machine mains power to a usable input to the DC power supply. When this supporting infrastructure can be eliminated control system DC power infrastructure is simplified and costs become lower. As a result, the SPE standard includes the provision for optional power delivery.

## THE SPE POWER ADVANTAGE

It is the provision for remote powering of devices connected with SPE that is the most significant and transformative aspect of this technology for industrial applications. The IEEE 802.3bu standard provides for the option of remote DC power over the SPE connection. This technology, Power over Data Line, or PoDL (pronounced "poodle"), is akin to the well-known and well-adopted Power-over-Ethernet (PoE) technology for standard Ethernet. The significant innovation here though is that both the data transmission and the power delivery is over the same pair of conductors, unlike PoE which requires a separate pair for power delivery to the device. This is why PoE technology can't be used on SPE. PoDL represents a necessary adaptation of PoE for SPE.

SPE and PoDL work together as technologies to provide both simultaneous data transmission and device power over the same link. Also, PoDL can be used to power a device with no data being transmitted. Like PoE, PoDL includes communication with devices to determine appropriate power levels. The 802.3bu standard defines multiple voltages and power classes. Currently, the standard supports power levels up to 1.6 amps per conductor that can be delivered to devices at the end of a 1km SPE link. Provisions for power levels up to 2 amps per conductor is expected soon.

This technology means that for many devices, local DC control power and its associated infrastructure will not be necessary, which is truly transformative for the manufacturing and industrial network edge.

SPE, together with PoDL, provides an elegant solution for connecting field devices to the factory and plant-wide network. Protocol translators and gateways are eliminated, greatly simplifying the overall network topology and providing one that is easier to maintain. In addition, with remote

**The advent of SPE will significantly evolve the industrial network edge, paving the way for a seamless network fabric that will be the foundation of Industry 4.0 and the IIoT.**

power to the device provided by the network switch or injector, control power infrastructure and provisioning is not needed. Further, as compared to multi-pair Ethernet, the single-pair Ethernet PHY chip electronics are much simpler by design. This, along with much simpler cable and connector designs, will all result in a network that is more reliable.

**PERSPECTIVES & CONCLUSION**

SPE is a forthcoming transformative technology that will soon be commercially available to manufacturers and industrial plant facilities. For the manufacturing space, a precedent, or model for success has been set by the automakers and their ecosystem of suppliers with the development and adoption of SPE (100BASE-T1) for in-vehicle networks. Other market segments such as rail transportation and building automation will also begin to adopt SPE networks in the coming years.

SPE as an enabling technology for manufacturers and industrial plants to:

- Cost-effectively achieve a single, seamless network from cloud to edge
- Supply edge devices with data and power on a single connection
- Gain significant bandwidth at the edge versus legacy protocols
- Simplify edge networks by eliminating protocol translation gateways
- Transform and simplify DC control power infrastructure
- Improve cybersecurity by eliminating legacy protocols
- Connect miniaturized micro-IoT and otherwise constrained form-factor devices
- Connect field instrumentation in hazardous environments with Advanced Physical Layer (APL, for intrinsic safety)
- Lower total cost of ownership for the OT network

Progressing into late 2020 and early 2021, look for product roadmaps and offerings from device, connectivity and network equipment manufacturers. Reference architectures, like Converged Plantwide Ethernet (CPwE), will be updated to incorporate SPE and APL into validated designs with vertical application focus.

At the Rockwell Automation's Automation Fair held in Chicago during November 2019, Panduit demonstrated a functional SPE connection between a human-machine interface and a device over a kilometer (1km) of cable consisting of 10 in-line connections. The demonstration generated much excitement and interest from many who witnessed it. Other industry solution providers at the event also demonstrated SPE technologies.

In short, the ecosystem is actively developing SPE solutions that are attractive to the manufacturing and industrial marketplace. The advent of SPE will significantly evolve the manufacturing and industrial network edge, paving the way for a seamless, "cloud-to-edge" network fabric. This will create the foundation for Industry 4.0 and IIoT, enabling both quantity and time precision of information for enterprises to gain more control over and insights into their underlying processes. ◻

# Edge, cloud or something in between?

**By Bob Sperber**

Analysts predict the number of semiconductors shipped for use in IoT applications overall will grow from 30-something billion chips today to more than 70 billion annually in the next five years. Meanwhile, enormous cloud datacenter campuses continue to proliferate worldwide, and are measured on an energy consumption scale once reserved for the power plants that feed them.

On the edge, in the cloud and embedded within the pervasive networks that weave these digital worlds together, the designer of industrial automation and information management solutions today has at his or her disposal computing power that was unimaginable only a few years ago.

Increasingly, it's not a question of computing resource availability, but rather where it's best to solve which types of problems. Oftentimes the edge wins out for reasons of performance and determinism. Other times, the cloud carries the day because new applications can be spun up more quickly and readily address problems that are broader in scope than a single unit or facility.

Indeed, an informal tour of industry leaders indicates that the edge or cloud seldom takes all. Rather, an integrated, hybrid architecture that features a combination of edge and cloud execution promise the best of both worlds.

## CHIPMAKER, HEAL THYSELF

Take Intel, for example. The semiconductor manufacturer is both creator and user of the technology that makes the digital industrial age possible, and it recently completed its first test of a scaleable, edge-to-cloud predictive maintenance solution for use in its own fabs.

The project sought to automate the labor-intensive process of monitoring fan filter units (FFUs) that purify air for manufacturing. To optimize vibration data collection from each FFU and allow quick event triggering, Intel leveraged its GE Digital factory automation platform and its own Intel IoT gateways to create an edge solution. Locating data processing on-site made more sense in terms of processing power, cost and time.

"The fact that we analyze at the edge will optimize the traffic that crosses the network. If we were to send all data to the cloud, the solution would be far too expensive," says Chet Hullum, Intel's general manager, Industrial Internet of Things. Yet there remains a role for the cloud: collecting summary-level data for long-term trend analysis. The project has been a success, based on results including:

- A more-than 97% increase in FFU uptime due to early parts orders upon potential failure detection, and rapid replacement.
- 300% reduction in unscheduled downtime over manual, labor-intensive FFU monitoring, which in turn boosted productivity.
- An estimated reduction in cloud traffic of nearly 94%.

**Increasingly, it's not a question of computing resource availability, but where it's best to solve which types of problems.**


Intel uses edge analytics together with cloud-based trend analysis to head off failures of fan-filter units in its semiconductor fabs. (Image courtesy of Intel Corp.)

Hullum calls the foray into edge computing an "ideal example" of leveraging an IoT gateway at the "point of ingestion" and optimizing system scalability. He reports that Intel is now expanding the FFU solution to more production lines. The company is also expanding the edge-based predictive maintenance architecture "to two additional use cases this year," including one project to predict pump failures before they occur.

This sort of solution squares neatly with the Hewlett Packard Enterprise (HPE) view of an applications architecture that is "edge-centric, cloud-enabled and data-driven," said CEO Antonio Neri, in a keynote address to attendees of the ABB Customer World 2019 event in Houston.

HPE has long contended that the new digital architecture for industry will be a hybrid, edge-centric one due in part to compliance (data ownership/privacy), latency and bandwidth issues. "75% of all data is created at the edge, where we live and work," Neri said. Meanwhile, only 6% of that data is put to use, and sending it to the cloud can make matters worse, Neri added, likening the cloud to The Eagles' Hotel California**:** "Once it's in, it can be really hard to check your data out."

"We believe the better solution is an edge-to-cloud architecture, where you only move data to the cloud as needed," Neri said. "It's all about managing that data effectively, and extracting outcomes faster.

Manufacturers and other industrial organizations have been collecting and analyzing data for decades—historically in centralized data centers that can be quite distant from the factory floor, notes Bob Voss, senior principal engineer at Panduit's Jack E. Caveny Innovation Center, the company's corporate R&D center in the southern suburbs of Chicago.

"Many newer industrial automation applications require real-time

or near real-time interaction with compute resources," Voss says. "Due to this requirement, many factories are moving the compute resources closer to—or on to—the factory floor. This opens up a whole other realm of issues when you expose datacenter servers and networking equipment to these harsh environments. Contamination from water, dust and corrosives become potential hazards. Appropriately protective enclosures, equipment and physical infrastructure are required to properly address these issues."


"It's all about managing that data effectively, and extracting outcomes faster." HPE CEO Antonio Neri subscribes to an "edge-centric, cloud-enabled, data-driven" approach that moves data to the cloud "only when needed."

## NEW REVENUE IN NEW SERVICES

Industrial equipment manufacturers also are developing new cloud services enabled by local, edge-based analytics for assets in the field. Caterpillar, for one, recently built on cloud services from OSIsoft to offer its own Asset Intelligence platform to analyze fuel consumption, equipment health and other critical operations. These cloud-based services, which work in conjunction with analytics performed locally, helped one operator of large marine vessels save $450,000 in fuel per ship annually by optimizing hull-cleaning maintenance to reduce drag. The service also helped saved a cruise line $1.5 million per ship in reduced fuel consumption.

Elsewhere, Gardner Denver, global provider of industrial equipment, employed Software AG's Cumulocity IoT platform to launch a subscription-based condition monitoring service for its IoT-enabled compressors. Users can remotely monitor operational parameters in near real-time, and receive notice of fault conditions. These instances of remote monitoring, maintenance and management of manufactured assets show how remote edge analytics can be parlayed into cloud-services that bring new value to stakeholders.

Industry players also are looking for the killer use case to unlock new analytics-driven services. One new arena that's got Steve Carlini excited is electrical energy storage and grid management. As the traditional, centralized energy grid continues shifting to distributed generation and micro-grids, lithium ion battery banks will begin to emerge as a key resource for buffering supply and demand, says the Schneider Electric vice president of innovation and data center. "With the right amount of information and the right analytics, you can start discharging these batteries to cut electric costs or supplement power when it's needed."

## EDGE ANALYSIS + CLOUD PERSPECTIVE

Industry leaders also are seeking to push analytics to the next level: from predictive to prescriptive. Currently, 50% of industrial firms

Mitsubishi Electric Automation has alternately put edge and cloud strategies to work on a range of manufacturing challenges. One use case for Oracle cloud applications is to increase the accuracy of pick-and-place robots, improving the visual detection and rejection of off-spec product. (Image courtesy of Oracle.)

have, or are piloting, an industrial analytics program, while another 48% of companies plan to within the next three years, according to LNS Research. Firms are, however, finding it difficult to break into the prescriptive realm because "it takes far more information than is available at the edge," says Dan Miklovic, LNS Research fellow. There's new value to be found in systems that can tell operators, "run this bearing at X speed to meet the production schedule, then take it offline for service," Miklovic says. Now, edge-to-cloud systems are targeting such solutions.

NRG Energy, which supplies electricity to more than 38 million U.S. households, credits prescriptive analytics with increasing turbine efficiency to save an estimated $5 million a year, with zero impact on planned outage schedules. These results are based on the company's implementation of GE Digital's Predix platform to bring essential turbine data from the edge to the cloud. There, analytics blend real-time production data, external lifing models for turbine components, and periodic pricing and weather reports "to let plant operators know the most profitable way operate the turbine," explains Amy Aragones, senior director of product management, GE Digital.

This hybrid edge/cloud solution analyzes how and when to run turbines safely beyond baseload conditions during periods of peak market prices, and then guides users how and when to under-fire the turbines to recoup the wear incurred during the peak-profitability hustle. This reportedly preserves both expected turbine service life and planned maintenance schedules.

For leading photovoltaic manufacturer First Solar, most analytics are performed in the cloud. Beyond the on-premise capabilities of its Rockwell Automation control and information management infrastructure, data from virtually all machines, PLCs, robots and other

IoT-enabled devices are sent to the cloud for deeper and more broadly based analysis. In one month at First Solar, four plants send five billion manufacturing database records to the cloud, each containing approximately 100 data points on equipment performance.

Cloud-based services makes sense because "We're not looking at a single piece of equipment at a single location, we're looking to compare all similar equipment at every one of our locations," says Allen Blackmore, IT domain architect for global enterprise technology, First Solar. He envisions analytics on an enterprise-wide data lake of unstructured data across manufacturing, sales, finance, and supply chain functions—essentially the entire business—which by its scope necessarily transcends the limits of an edge-only approach.

## MODERNIZING MADE PROFITABLE

Bringing IoT capability to legacy assets "is one of the biggest issues that enterprises face," says Ricardo Buranello, vice president of global factory solutions at Telit, an IoT infrastructure provider. But it's worth the effort, he says, citing results achieved at a customer that manufactures automotive axles. By connecting and analyzing the data from 1,000 formerly isolated CNC machines, the company was able to increase productivity and pocket savings worth an additional 50 assemblies per day.

"We're seeing an enormous increase in the tag values companies would like to collect from their legacy machines, like 30 year-old lathes," says Dave Cronberger, infrastructure architect with Cisco Systems, adding that manufacturers are showing "a strong belief that they're going to gain new insights from their controls and I/O blocks, and learn new things that they don't know currently."

Mitsubishi Electric Automation has alternately put edge and cloud strategies to work on a range of manufacturing challenges. When the use case called for real-time analytics to predict and improve electroplating quality or to control defects in injection molding, they focused on local execution: "It would have cost us a fortune to put all our production data into the cloud; we'd run out of space in 10 minutes," says Timothy Lomax, strategic alliance manager.

Elsewhere, Mitsubishi has used Oracle's cloud and business applications to increase the accuracy of pick-and-place robots, improving the visual detection and rejection of off-spec product. The project, now underway, uses the cloud for robot data, images, data trending and artificial intelligence analyses. For many projects at Mitsubishi and elsewhere, Eric Prevost, vice president and global head of emerging technologies for Industry 4.0 at Oracle, reports working "in coordination with edge-device providers for many projects."

Research from Wikibon predicts the evolution of a coordinated edge/cloud network model from sensor to supply chain. The research firm modeled a small wind-farm 200 miles from the cloud data center with IoT-connected security cameras, security sensors, sensors on the wind-turbines and access sensors for all employee physical access points. The result: When an edge network handled 95% of the data traffic for video and sensors, total cost was "reduced from about $81,000 to $29,000 over three years," about one-third the cost of the cloud-only approach.

Given the caveat that every organization has unique conditions, the mainstream belief that the "you can't do everything in the cloud" rings true, says Jason Andersen, vice president of business line management for Stratus Technologies, maker of high availability edge-computing solutions. "Today, from what we understand, the breakeven point seems to be around 30%," Andersen says. "Processing 30% of your data in the cloud will cost about the same as doing it all at the edge." □

# MQTT getting what it needs to go industrial

Lightweight messaging protocol is now the most commonly used above HTTP

**By Ian Verhappen, Senior Project Manager of Automation, CIMA+**

Message queuing telemetry transport (MQTT), developed in 1999, is a publish/subscribe message lightweight protocol based on TCP that is now the most commonly used messaging protocol above HTTP. The reference architecture is very simple, and is based on client/server. The client is generally a sensor that "publishes" the information to the server (broker) that receives the information and dispatches it to the subscribers. MQTT protocol uses a many-to-many paradigm, and the broker decouples the publisher to the subscriber and acts as a message router with every message a discrete chunk of data, opaque to the broker. MQTT's publisher/subscriber model enables clients to communicate one-to-one, one-to-many and many-to-one.

Every message is published to an address, known as a topic. Clients may subscribe to multiple topics. Every client subscribed to a topic receives every message published to the topic. The MQTT specification does not dictate any particular Topic Namespace, nor does it dictate any particular payload data encoding, though MQTT topics are hierarchical, like a filing system (e.g. sales volume/flow/corrected). Wildcards are allowed when registering a subscription (but not when publishing), thus allowing whole hierarchies to be observed by clients.

Only the most recent persistent message is stored. When a client subscribes to a topic, any persistent message will be sent to the client. However, unlike a message queue, MQTT brokers do not allow persistent messages to back up inside the server.

**MQTT's publisher/subscriber model enables clients to communicate one-to-one, one-to-many and many-to-one.**

MQTT also supports three quality of service levels: "fire and forget," "delivered at least once," and "delivered exactly once." To prevent excess traffic when a device knowingly disconnects, MQTT clients can register a custom "last will and testament" message to be sent by the broker if they disconnect.

MQTT also has support for persistent messages stored on the broker. When publishing messages, clients may request that the broker preserve the message.

MQTT brokers can require username and password authentication from clients to connect, and to ensure privacy, the TCP connection may be encrypted with SSL/TLS.

As a machine-to-machine (M2M)-oriented protocol, MQTT is designed to be lightweight, and it has two drawbacks for very constrained devices:

- Every MQTT client must support TCP and therefore, always holds a connection open to the broker, which can be a problem

**Sparkplug provides an open specification for how edge-of-network devices and applications communicate bidirectionally within an MQTT infrastructure.**

in high-packet-loss/computing environment.
- MQTT topic names are often long strings, which make them impractical for 802.15.4 industrial wireless environments.

These shortcomings are addressed by the MQTT-SN protocol, which defines a UDP mapping of MQTT and adds broker support for indexing topic names.

MQTT is an OSI Application Layer (Layer 7) like HTTP, and as a carrier, requires tagging such as HTML or XML for web pages to represent the data. MQTT for SCADA applications have Sparkplug, which provides an open specification for how edge-of-network (EoN) gateways or native MQTT-enabled end devices and MQTT applications communicate bidirectionally within an MQTT Infrastructure.

The Sparkplug specification defines an MQTT Topic Namespace, payload, and session state management that can be applied to the SCADA/IIoT market sector. To meet these requirements, the specification is based on a lightweight, bandwidth-efficient, low-latency payload encoding architecture.

MQTT-enabled infrastructure requires that one or more MQTT servers be present in the infrastructure. Typically, there will be only one primary SCADA/IIoT host node responsible for monitoring and control of a given group of MQTT EoN nodes. The SCADA/IIoT host node is any MQTT client application that subscribes to and publishes messages defined in this document.

The Sparkplug Topic Namespace define nine message-type elements:

Using these defined messages host SCADA/IIoT applications can:
- Discover all metadata and monitor state of any EoN/Device connected to the MQTT infrastructure.
- Discover all metrics, which include all diagnostics, properties, metadata and current state values.
- Issue write/command messages to any EoN/Device metric.

The above standards are all available and open, with Sparkplug from Cirrus Link Solutions, while Eclipse Foundation has released an open-source implementation of MQTT called Mosquitto, while OASIS maintains Advanced Message Queue Protocol (AMQP). Do they define a fully open solution? Not quite. But I have also heard rumors of an end user-driven group similar to the Open Process Automation Forum being discussed to support an Open SCADA standard.

## REFERENCES

1. Sparkplug MQTT Topic & Payload Definition Version 2.1, Cirrus Link Solutions, April 2019.
2. MQTT Version 3.1.1 Committee Specification, OASIS, April 2014.