



Edge Computing in Industrial Automation

IIoT Intelligence Starts at the Edge



CBC
CARLTON-BATES COMPANY
A SUBSIDIARY OF WESCO DISTRIBUTION, INC.

OPTO 22
Automation made simple.

Opto 22

43044 Business Park Drive • Temecula • CA 92590-3614

Phone: 800-321-6786 or 951-695-3000

Pre-sales Engineering is free.

Product Support is free.

www.opto22.com

Form 2184-180208

© 2016–2018 Opto 22. All rights reserved. Dimensions and specifications are subject to change. Brand or product names used herein are trademarks or registered trademarks of their respective companies or organizations.

EDGE COMPUTING IN INDUSTRIAL AUTOMATION

IloT Intelligence Starts at the Edge

On the train to work, Lee opened an email on her smartphone sent from a controller operating a surface-mount tool at her factory. Attached to the email was a quality control report that suggested changing the tool's solder temperature.

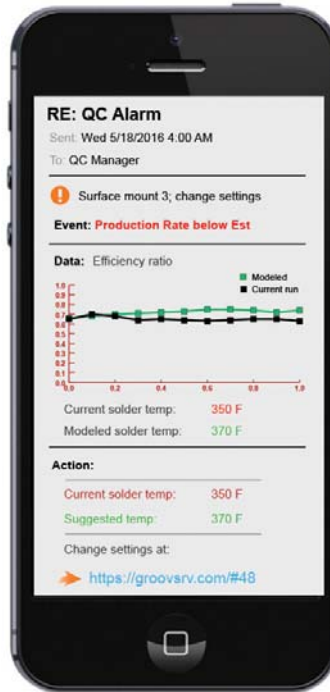
To generate that email suggestion, the controller had securely sent yesterday's production data to a cloud-based analytics system to compare current and historical data for the machine.

Next, it accessed the machine manufacturer's website and obtained the latest recommended settings.

Finally, it built a production efficiency report with a suggested solder temperature for today's production run that would increase yield by 7 percent over yesterday's run.

Lee clicked a link in the email and connected to the controller's mobile operator interface over a secure, encrypted channel. Lee logged in and navigated to the machine's solder temperature setpoint, where she entered the recommended value.

All this took place before she got to the office.



AT THE EDGE

That controller operating the surface-mount tool at Lee's factory operates at the edge of the factory's network.

Systems like these at the network edge are increasingly able to leverage cloud-based resources to perform edge computing—if computing resources exist as needed along the path from a sensor to the cloud—and if these computing resources reduce the total amount of data to be sent to the cloud for storage, processing, and analysis.

As a result, businesses can more quickly identify real opportunities for operational efficiency improvement and meaningful revenue generation.

To foster such business benefits, data from the physical world of machines and equipment must be available to the digital world of the internet and information technology systems, quickly, easily, and continuously.

Successful industrial internet of things (IIoT) applications require operational technology (OT) professionals to make data from their systems, which monitor and control the physical world, accessible to the data processing systems of information technology (IT) professionals.

Once the data is there, cognitive prognostics algorithms running on IT systems can analyze it, refining raw physical data into actionable information that can predict outcomes in real time.

The results can be used to improve inventory management and predictive maintenance and reduce asset downtime.

But before such benefits can be realized, three problems need to be solved: connectivity, big data, and IIoT architecture. Lee's factory has a new kind of controller, a *groov*® edge programmable industrial controller (EPIC),

Edge Computing in Industrial Automation

which goes a long way toward solving these three problems.

THE CONNECTIVITY PROBLEM

The internet of things runs on vast amounts of data, generated by the physical world and then transported and analyzed by the digital world.

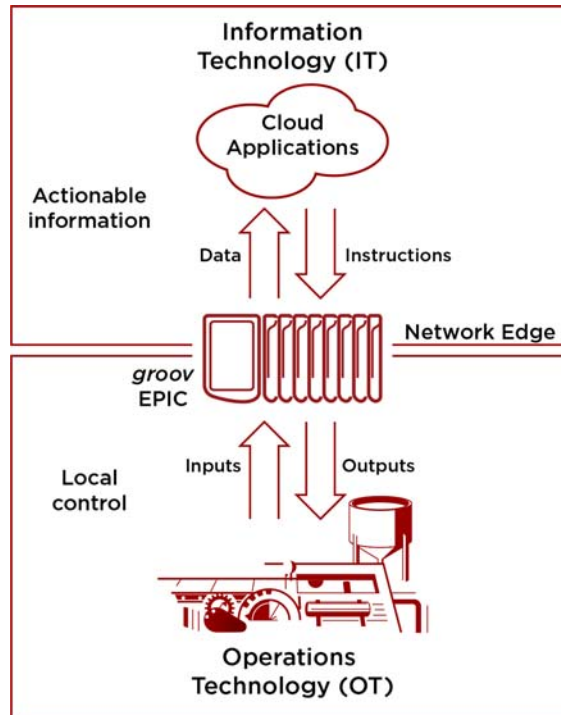
It's an attempt to achieve perpetual connectivity and communication between people and things and even between things and other things.

But in the industrial world, most of these things were never designed to serve this new purpose. They were designed and installed long before the internet was developed.

At the edge, things like sensors, circuits, relays, and meters are attached to industrial control systems used to operate equipment and machines. These sensors translate what's physically happening in the world (temperature, light, vibration, sound, motion, flow rate, and so on) into an electrical signal like voltage or current.

The voltage or current is then interpreted by a controller, which monitors and controls physical equipment and machines.

Sensors typically have little or no intelligence. They are designed to merely observe and report. They can't speak in the digital bits and bytes, the ones and zeros that information technology and internet devices understand and use to communicate.



They also lack the physical connections and logical interfaces to communicate on the industrial internet of things (IIoT). They do not have a built-in Ethernet jack or a wireless interface.

They don't understand the languages the internet uses, like JSON, RESTful APIs, and JavaScript. They don't run an operating system or have a built-in TCP/IP stack or web server.

And sensors have little or no built-in computing power, so providing edge computing at the sensor level to filter volumes of data before forwarding to the cloud is impossible.

So right now the internet and the industrial things we want to connect to it aren't communicating.

What can we do?

One option is to simply wait for highly intelligent, connected sensors to become available to the marketplace. But those sensors are years away from being cost effective.

Moreover, sensors installed today or even decades ago are still performing their tasks. They're just not connected to the IIoT, so the data they generate is siloed and inaccessible to IT systems for further analysis.

In Lee's factory, the *groov EPIC*® controller takes on the task of communicating with IT systems and the internet. Located at the edge and connected to sensors and devices

What is Edge Computing?

You've heard about cloud computing, which is using a network of remote servers, rather than a local server or personal computer, to store and manage data and run computer programs.

The advantages of cloud computing are that people or organizations can combine and share computer resources rather than having to build and maintain them.

Edge computing basically brings cloud computing down to the edge of the network, in the physical world. Edge computing is using computing power at the edge to filter or process data and then send only the required data to the cloud.

The advantages of edge computing are many. First, it reduces traffic on networks and the internet by reducing the amount of data sent. It also plays a valuable role in efficiency, security, and compliance.

Right now the internet and the industrial things we want to connect to it aren't communicating. There's a disconnect between the physical world of current and voltage and the digital world of servers and clouds.

through its I/O, the EPIC can speak the languages computer networks understand.

THE BIG DATA PROBLEM

Across the globe a massive installed base of things exists today, generating useful data that the IIoT wants to access and consume. In oil and gas applications, a typical oilfield has up to 30,000 sensors installed. Factories and plants across the world have billions of sensors.

Each sensor is capable of generating huge amounts of data from the physical world. Some IIoT applications could potentially generate terabytes of data per second.

These are volumes of data the digital world has never seen before. This is the Big Data problem.

Moving that much data onto existing network and internet infrastructures for cloud-based analytics and centralized management will clog networks, vastly increasing network and internet latency. For many industrial IoT applications, that is not acceptable, because real-time control and monitoring are mandatory.

For the industrial internet of things to reach critical mass, intelligence must be pushed to the network edge, where the physical world meets the digital world. Computing systems at the network edge must have the capability to collect, filter, and process data generated at the source, before it's transmitted up to the IIoT.

And at the same time these edge computing systems must be able to complete the local real-time process control and automation tasks of traditional industrial applications.

The Big Data problem can be at least partly solved by edge systems like the EPIC controller at Lee's factory, which has the computing power to sift and process data before sending it into network and internet infrastructure. With a real-time, open-source operating system, an industrial quad-core processor, and solid-state drives, a controller like this provides the needed intelligence. And because it is first

a programmable industrial controller, the EPIC also provides local monitoring and control.

THE IIOT ARCHITECTURE PROBLEM

Our third problem revolves around how today's IIoT architecture works. Let's take a look at its complexity and explore a possible path forward.

For a cloud-based server to capture data from an analog sensor today, the sensor's data must be translated using a series of software and hardware tools.

First, the sensor is physically wired to a device such as a PLC (programmable logic controller). While modern PLCs do provide basic analog-to-digital conversion of sensor signals, PLCs were not designed to interface with the IIoT.

PLC hardware, software, and programming languages were designed for repetitive, application-specific tasks like process control and discrete automation. They typically use proprietary protocols and languages for communication and programming, and do not include information security standards like encryption and authentication.

PLCs were originally designed as standalone systems. The protocols they use are seldom internet compliant and are designed for point-to-point communication instead of the point-to-multipoint communication architecture found in the IIoT ecosystem.

If systems that communicate using internet-compliant protocols—such as PCs, web servers, and databases—want to communicate with a PLC, a vendor-specific and

What is an Input/Output (I/O) System?

Things in the physical world typically communicate in an electrical signal like voltage or current, but the digital world doesn't understand these signals. I/O systems act as the translator.

I/O systems are wired directly to sensors and actuators. Their job is to convert analog data (like voltage or current) from a sensor into digital data (ones and zeros) to send to a computer, or to convert digital data into analog data to send to an actuator, to take action in the physical world.

An input sends data from the physical thing to the computer; an output does the reverse. Inputs and outputs are like the fingers on a hand, sensing the world, sending data about it to the brain, and then taking commands from the brain back to the fingers to take action.

Edge Computing in Industrial Automation

often proprietary software driver or hardware-based protocol gateway is required.

OPC (Open Platform Communication) software is one solution to this communication disconnect. But OPC was originally designed around PC architecture using the Microsoft® Windows®-only process exchange, COM/DCOM. Most systems and devices connecting to the IoT are not Windows-based devices.

For example, take your smartphone. It's likely an Apple® or Android® device, both of which run modified versions of the Linux operating system, where COM/DCOM process exchange does not exist.

OPC UA (Unified Architecture) has been released, but it's merely a wrapper for existing OPC drivers built on Windows architecture. It requires design engineers to build an OPC UA client adapter into their products. And even then, modern network and internet assets such as web servers, databases, smartphones, and tablets do not speak OPC UA.

PLCs, OPC servers, proprietary drivers, and protocol gateways quickly become a convoluted IIoT architecture. These layers of complexity not only require time, money, and specific domain expertise to install and maintain, but also the data being sent from the physical world has been converted by so many different pieces of hardware and software that data integrity can be jeopardized.

Imagine the difficulty in provisioning and troubleshooting these IIoT systems.

And then consider that today's automation architectures often do not address information security. Sending data generated at the edge through so many layers of conversion opens up complex information security concerns as the data is transported to the cloud.

Multiply these architectural issues across the billions of devices we expect to connect using the IIoT, and you see the communication challenge the IIoT faces.

Flattening the IIoT Architecture


One way to simplify this complex IIoT architecture is to consider a different communication model.

Most devices on a computer network today communicate using the **request-response** model. A client requests data or services from a server, and the server responds to the request by providing the data or service. This is the model company computer networks use, and it's also the typical model for that biggest of all networks, the internet.


Each client on the network opens a direct connection to each server and makes a direct request. A client doesn't know when the data changes, so it sends requests at regular intervals (in automation, as fast as once per millisecond), and the server responds each time.

This communication model works very well on a network where the number of clients and servers is not large. But the IIoT poses problems of scale. If each client must be connected to each server it uses and must constantly request and receive data, network traffic can quickly become an issue.

The Problem: Complex Current IIoT Architecture



The Solution: Flattened IIoT Architecture



Another model may be more effective in IIoT applications: **publish-subscribe**, or pub-sub.

In the pub-sub model, a central broker becomes the clearinghouse for all data. Because each client makes a single lightweight connection to the broker, multiple connections are not necessary. Clients can publish data to the broker, subscribe to data going to the broker from other clients, or both.

In addition, clients publish data only when it changes (sometimes called *report by exception*), and the broker sends data to subscribers only when it changes. The broker does not hold data, but immediately forwards it to subscribers when it arrives. So multiple data requests are not necessary.

These two factors—just one connection per client and data traveling only when it changes—shrink network traffic substantially, making pub-sub an effective communication model if you have many clients and many servers.

Pub-sub can also be effective when it's difficult to set up a direct connection between a client and a server, or when the network is low-bandwidth or unreliable—for example, when monitoring equipment in remote locations.

The pub-sub transport protocol MQTT is frequently mentioned for IIoT applications, especially when used with Sparkplug messaging, which was designed for industrial applications. Secure data communications using MQTT/Sparkplug can help flatten IIoT architecture and more easily produce the results we need.

ARCHITECTURAL CHANGES

As we've seen, for the IIoT to reach critical mass, internet protocols and technologies need to be driven into systems at the edge, where the physical world and the digital world connect.

Layers of complexity must be removed from the communication process between digital systems and physical assets. Modern IIoT system architectures must be flattened, streamlined, optimized, and secured.

Edge computing systems must easily and securely access the cloud through the open, standards-based communication technologies the internet is based on.

That means:

In the long run, OT/IT convergence will demand a flattened architecture and seamless communication between assets, using open, standards-based communication protocols and programming languages.

- **Internet technologies** like MQTT/Sparkplug, TCP/IP, HTTP/S, and RESTful APIs—the dialect of the internet—must be built directly into devices at the edge.
- Internet **security** technologies like SSL/TLS encryption and authentication must be built in directly to edge computing systems.
- Cloud-based systems must be able to make **RESTful API** calls to access data, or use a **publish-subscribe** communication model like MQTT/Sparkplug to get data from remote edge devices, without the layers of complexity and conversions that exist in industrial applications today.

THE POWER OF INTEROPERABILITY

The standard technologies used by the internet for transmitting information have created a cohesive system for communication. But we have not always had this cohesive system.

Before the internet and the world wide web, many different internet-like protocols and architectures existed. Computer systems all ran different operating systems, requiring different programming languages.

Small pockets of interconnectivity existed, but for the most part systems were disconnected from each other. It was very similar to the way industrial systems communicate today, with the need for converters, adapters, and gateways.

The internet was designed to allow input/output and information systems to share data through a common interface, removing layers of complexity and allowing for greater interoperability between systems designed and manufactured by different vendors.

That's why an Apple computer or Android phone today can send an email to a Windows computer: despite their incompatibilities, they speak the same internet languages. Today's internet uses a common set of protocols, tools, and

Edge Computing in Industrial Automation

routines designed to make the transportation, acquisition, and analysis of digital information a seamless process, no matter what device you're using.

Although sensors and other physical assets installed at the edge may not have been designed with internet interoperability in mind, there's still a massive opportunity to collect meaningful data from the huge installed base of existing things.

But it will require a solution that understands both sides of the OT and IT convergence—something that can:

- Locally **translate the physical world** of currents and voltages (OT) into the secure communication protocols and languages the digital world (IT) understands.
- **Process and filter mountains of data**, sending only the necessary data to the cloud for analysis.
- Provide communications interfaces and processing power to **maintain the closed-loop, real-time control requirements** of industrial applications.
- Deliver all of the above in a package **suitable for challenging industrial environments** where dust, moisture, vibration, electro-mechanical frequencies, and temperature vary widely.

CONCLUSION

We've seen that edge computing is the sensor on-ramp to the IIoT. Once the communication, security, and computing technologies of the internet find their way into computing at the edge, the IIoT will begin to reach its potential.

internet technologies are available in some industrial systems today. And some vendors have already started bridging the gap between OT and IT by adding IIoT technology like MQTT and RESTful APIs directly into their controllers.

Our shortest path to a successful IIoT is to leverage the existing interoperability technologies of the internet in industrial automation products and applications.

HOW CAN OPTO 22 HELP YOU?

At Opto 22, our goal is to flatten the IIoT architecture so you can realize your IIoT goals.

Opto 22's engineering focus is on building hardware and software tools to bring the benefits of the IIoT—simply, reliably, and securely—to the things that already exist in your world.

Our products offer an easy and cost-effective way to bridge the real world with the digital world.

- Reliable Opto 22 input and output (I/O) modules connect with virtually any electrical, electronic, mechanical, or environmental device, converting these raw signals to useful digital data.
- Opto 22's *groov* EPIC controller offers solid industrial control, logic solving, data collection and processing, and the ability to move data over the standard networks and protocols understood by IT.
- When it comes time to visualize, notify, and mobilize your information, *groov* View (included in the EPIC controller) offers a simple way to build mobile operator interfaces that can be used on any screen, from your smartphone to big-screen HDTV.



Opto 22's *groov* EPIC edge programmable industrial controller

groov EPIC connects directly to sensors and actuators at the edge, monitors and controls the physical world, processes data locally, and communicates that data using internet-standard protocols and languages.

All Opto 22 products are backed by more than four decades of expertise in applications like process control, discrete manufacturing, remote telemetry, data acquisition, and supervisory control.

Opto 22 products are supported by our experienced engineers at no charge and available worldwide.

ABOUT OPTO 22

Opto 22 was started in 1974 by a co-inventor of the solid-state relay (SSR), who discovered a way to make SSRs more reliable.

For over 40 years, we've brought commercial, off-the-shelf technologies to industrial systems all over the world, designing our products on open standards. We pioneered the use of PCs in controls back in the 1980s, Ethernet networking at the I/O level in the 1990s, and machine-to-machine connectivity in the 2000s. Today, we build secure internet technologies into our controllers and I/O.

All Opto 22 products are manufactured and supported in the U.S.A. Because the company builds and tests its own products, most solid-state SSRs and I/O modules are guaranteed for life.

