Industrial Internet of Things and Communications at the Edge

by Tony Paine, CEO, Kepware Technologies

Contents

Summary

This eBook explores the Industrial Internet of Things (IIoT). It describes the benefits of Internet-enabling all hardware and software components ("Things") that comprise an automation system, and delves into the challenges the industry must overcome for the IIoT to be successful. These challenges include the ability to:

- Identify Things within the Internet
- Enable the discovery of Things and the data they possess
- Collect, move, and archive the massive amounts of data produced by the Things
- Handle intermittent to long-term outages in connectivity
- Integrate existing infrastructure into new IIoT strategies

But as the industry looks to build out the IIoT, its biggest challenge will be seamlessly enabling the Things that live on the edge of the network. The edge bridges the gap between IT and Operational Technology (OT), where the rich resources available in the cloud are not directly available. These industrial networks have their own challenges, including the use of:

- Disparate communications mediums
- Non-standard methods of identification on the network
- Different request/response models
- Short-term storage for small-scale, high-frequency data

After taking an iterative approach to working around these challenges, this eBook identifies the requirements and features of a new communications platform capable of seamlessly integrating industrial data into IIoT.

About the Author

Tony Paine is CEO of Kepware Technologies, a software development company focused on communications solutions for industrial control systems. His main focus is on the company's vision and long-term strategy around products and technology.

Tony has had a passion for integrating software and hardware since his early childhood, when he developed an application that turned a rudimentary text editor into a word processor with generic print capabilities. He focused his education around this interest and earned a Bachelor's of Science degree in Electrical Engineering, with a concentration in Computer Software and Hardware Design, from the University of Maine at Orono.

"While companies generate massive amounts of data, the insight that this data provides is only valuable if it can be accessed quickly and efficiently."

Tony joined Kepware Technologies as a Software Engineer in 1996 (shortly after the company was founded) and became instrumental in the design and implementation of the company's flagship product, KEPServerEX. He eventually took on the role of Chief Software Engineer where he led the development team, and was responsible for the creation of the company's LinkMaster and RedundancyMaster products. Throughout the years, Tony's role continued to evolve to reflect his growing responsibilities to the organization, first as Executive VP and Chief Technology Officer, and then as President and CEO in 2009. In 2014, Tony promoted Brett Austin to President in order to focus exclusively on long-term strategic planning as the company's CEO.

Tony has been involved in various technical working groups, where he has contributed to the direction and review of various standards used within the Automation industry. He currently sits on the Dean's Advisory Council for the University of Maine's College of Engineering, where he provides industry insight and evangelism around education in the area of technology.

1 Introduction

Introduction

IIoT Vision

The promise of connecting everything within an industrial environment to get complete visibility into its operations and allow the best real-time decisions to be made—with or without human intervention—will transform how we manufacture for years to come. The premise for this next industrial evolution is the Industrial Internet of Things (IIoT).

IIoT sets its sights on Internet-enabling all hardware and software components (the "Things") that comprise an automation system. This will allow for smarter automation where Things can share information, learn about their surroundings, and auto-tune themselves to achieve optimum throughput and minimal downtime. Personnel responsible for the operations of the system will be able to remotely assess and manipulate all aspects without the need for dedicated on-site expertise.

These benefits are contingent upon the resolution of key challenges—several of which the industry has been solving for years. In an industrial automation process, there are mechanical, digital, and human components. At any time, one of those parts may have information that is valuable to another part. In determining and adhering to technologies that connect those components, the Internet of Things has existed within industrial environments for some time now—just at a much smaller scale and under different names (like SCADA, M2M, Predictive Maintenance, and Process Optimization).

IIoT **Benefits**

Cost savings from:

- Minimized energy usage
- Integration with the supply chain
- Fewer on-site personnel
- Less time dedicated to low-level tasks

Better customer service through:

- Predictive maintenance
- Remote troubleshooting and patching
- Product improvement based on customer use

Smarter automation, providing:

- Auto-tuning and optimizing based on surroundings
- Notification reporting for diagnosis and resolution
- On-demand assembly driven by business systems

Visibility, any time and anywhere, enabling:

- Better decision-making
- Increased time to value
- Increased system safety and security
- Remote asset monitoring and managing

Today, there are several changes impacting the scale and speed of the IIoT. New vendors are entering the market, looking to consolidate data into actionable information, unify historical solutions, and bridge the gap between the public and private operational domains. Furthermore, our society is increasingly reliant on the Internet and has more connected tools available than ever before. Technology is no longer cost-prohibitive: we can network-enable anything with low-cost sensor technology, unlocking and storing data that was previously unavailable. Finally, the next generation of engineers are students growing up with technology that is rich, easy to use, and everywhere—creating an expectation that existing control systems will be comprised of technology that plugs in and works with little effort.

As industry looks to build out the IIoT, its biggest challenge will be seamlessly Internet-enabling the Things that live at the edge of the network.

At the enterprise level, multi-site awareness will provide critical insight for competitive strategic planning, as well as the opportunity to integrate beyond organizational boundaries for the purpose of leveraging a third party's business services.

As industry looks to build out the IIoT, its biggest challenge will be seamlessly Internet-enabling the Things that live at the edge of the network. Industry-wide, this area contains trillions of Things that contain one or many data points that may need to be analyzed and combined into information. Unfortunately, the edge of the network is also the furthest removed from the Information Technology (IT) we have become accustomed to using when Internet connectivity is required.

2

IIoT Challenges

IIoT Challenges

Identifying Things within the Internet

In order for Things to be able to communicate with each other, they need to be uniquely identifiable within the Internet. Historically, this has been accomplished through the assignment of an Internet Protocol (IP) address. As industry looks ahead to the trillions of Things that will be connected, focus has been placed on adopting the IPv6 standard, which defines a 128-bit address capable of uniquely identifying 340 undecillionth (340×10^{36}) addressable items (compared with only 4 billion addressable items using today's IPv4 standard). Though this range will more than cover the requirements of IIoT, it will be difficult—if not impossible—to manage this effectively on a global Internet scale. Typically managed by Naming and Number Authorities with the aid of Network Administrators, this will be an impediment as Things are added at an unprecedented rate.

Balancing ease of discovery with the rigid constraints of security will be fundamental to the success of IIoT.

Discovering Things and the Data They Possess

Once a Thing can be identified, the next challenge is how other interested parties will discover that it exists and what data it possesses. Of course, a Thing should be able to restrict discovery of all or some of its data based on security requirements. Balancing ease of discovery with the rigid constraints of security will be fundamental to the success of IIoT and must be achievable without a PhD degree in cybersecurity.

Managing Massive Amounts of Data

These trillions of Things will produce something much larger than trillions of data points (industry currently measures this in zettabytes or 10^{21} bytes), all of which will need to be collected, analyzed, and possibly archived. Moving this amount of data over the Internet will consume new levels of bandwidth, which could result in the degradation of service as well as higher costs for Internet carriers, service providers, and ultimately end-users. Moreover, archiving this data for future analysis will require massive amounts of data storage and a new generation of scalable applications for honing in on points of interest in a timely manner.

"ABI Research estimates that the volume of data captured by IoT-connected devices exceeded 200 exabytes in 2014. The annual total is forecast to grow seven-fold by the decade's end, surpassing 1,600 exabyte — or 1.6 zettabytes — in 2020."

Source: Business Wire

Navigating Connectivity Outages

The Things that make up IIoT, as well as the communication mediums that link them together, will not be available 100 percent of the time. While some downtime may be scheduled, there will be physical or environmental changes that result in intermittent to long-term outages. The severity will highly depend in cases where data loss is unacceptable or the criticality of variances in the data needs to be known in real-time.

Integrating Existing Infrastructure into New IIoT Strategies

Industrial Things have made data accessible over private networks for decades through the implementation of open or proprietary protocols. To achieve success in the areas of network optimization and third-party integration, complexities like security have largely been ignored. Given that the typical lifecycle for industrial Things exceeds twenty years, there will be an expectation to integrate the existing into new IIoT strategies. Opening these private networks and the data they contain to the Internet will require detailed security assessments to minimize risk of exploitation.

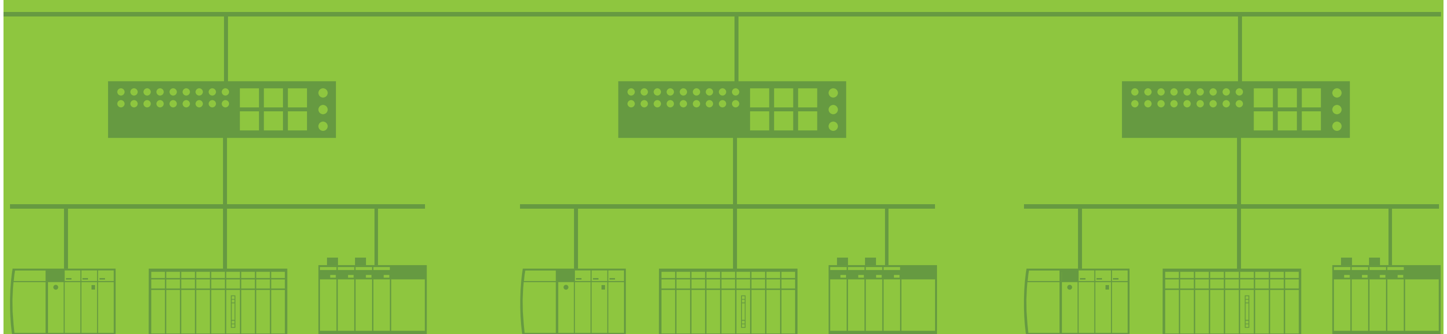

3 IIoT Edge

IIoT Edge

Leveraging the Power of Cloud Computing

In order to alleviate some of the preceding challenges, IIoT strategies will focus on pushing data into a centralized cloud platform. This platform and its corresponding services will be administered by IT experts familiar with the managed world of IP and made available to anyone with the proper credentials and an Internet connection. Leveraging the power of cloud computing and its multitudinous resources will make the required storage and processing power available to handle the zettabytes of data that will be collected, analyzed, and archived. Furthermore, the overall uptime of these platforms continues to trend upwards as they become more resilient to the increasing demand and expectations of our connected world.

Leveraging the power of cloud computing and its multitudinous resources will make the required storage and processing power available to handle the zettabytes of data that will be collected, analyzed, and archived.

Communicating with Devices on the Edge

The actual source of data pushed into the cloud resides within the industrial Things that live at the edge of the network. The edge bridges the gap between IT and Operational Technology (OT), where the rich resources available in the cloud are not directly available. OT encompasses industrial networks that have their own nuances and introduce additional challenges.

Connecting Disparate Communications Mediums

Very often, industrial networking technologies do not leverage Ethernet as their physical communications layer. Depending on the environment and the Things that comprise a system, anything from RS232/485 to modems to proprietary wiring may be encountered. Likewise, the data protocols that are exposed over these communication mediums are not likely to be IP derivatives. Consequently, a hodge-podge of industrial networks have been created with no attention made to the future possibility of being connected to the Internet.

Utilizing Non-Standard Methods of Identification

Unlike IP addresses in the IT world, many industrial Things have their own addressable schemes for uniquely identifying themselves on the network. These schemes vary by vendor and type, and may or may not have built-in discovery mechanisms. Innate knowledge by an integration expert is required to interconnect the Things in a way that makes them function as a whole.

Determining a Request/Response Model

Industrial Things have historically followed a request/response model. If a particular Thing is interested in a piece of data contained in another Thing, it will make an appropriate connection, request the piece of data, and wait for a response containing the result. Although this pull model is fine for Things living within the same digital boundary of OT, security and scalability requirements will dictate that this model is unacceptable for the outside IT world trying to look inbound. Instead, IIoT prefers a push model, where industrial data flows outbound to a cloud platform.

> *"IoT at a much smaller scale has been around within manufacturing for some time, just maybe under different names. . . The ability to connect the various pieces of equipment to computers that humans can interact with has been around for quite some time."*
>
> *Source: Manufacturing.net*

Enabling Short-term Data Storage

Within the context of a single industrial network, we may find up to thousands of Things that together may generate up to several thousands of data points. Though this sounds like a small set of data, the real-time requirements found in OT will require these points to be sampled at sub-millisecond rates for data change detection. In the past, this high frequency data would be simply analyzed, acted on accordingly, and thrown away. As we move to making this data available to IIoT, we will require short-term storage to ensure it can be pushed to other parties when they are available.

"Options are important, and sometimes the most effective manu-facturing infrastructure consists of products from various vendors. The difficulty lies not only in connecting those various brands, but also in ensuring that they can exchange information by supporting open protocols."

Source: InformationWeek

IIoT Edge Solution

4

IIoT Edge Solution

In order to seamlessly integrate industrial data into IIoT, a new communications platform is required. This platform requires extensive knowledge of the intricate realm of OT and the state-of-the-art and rapidly-changing domain of IT.

Within OT, the platform must understand the various network topologies and data protocols that will be encountered. It must be able to automatically discover and identify industrial Things and the data they contain, as well as be able to handle the storage of high-frequency updates.

Within IT, the platform must be able to transform the data it collects and push it into the cloud via IIoT standards. Emerging standards include Asynchronous Messaging Queuing Protocol (AMQP), Message Queueing Telemetry Transport (MQTT), Constrained Application Protocol (COAP), and Data Description Services (DDS). These standards allow for the retransmission of data in the event it does not reach its destination

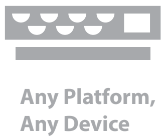

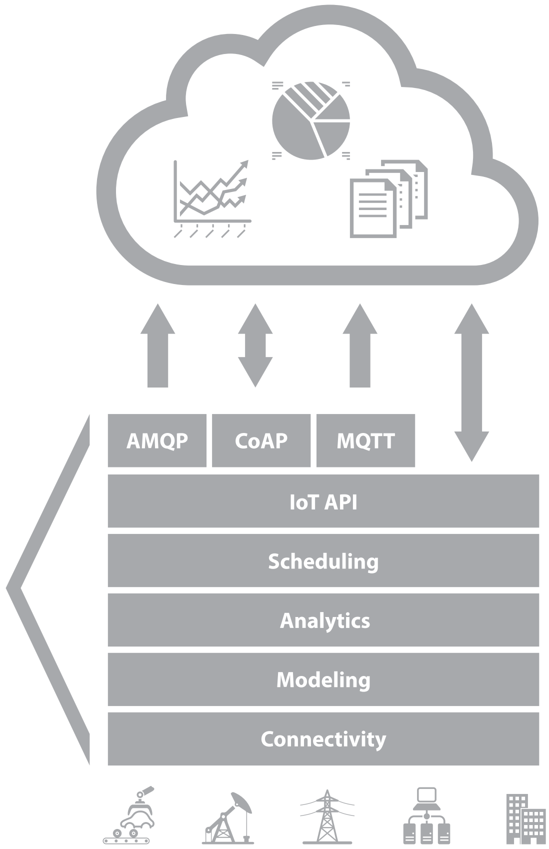

With the lack of computer networking infrastructure in OT, this platform must be embeddable and run within a stand-alone appliance or an edge-based switch or router where IT and OT converge.

Its flexibility will enable industrial data to be sampled cyclically or based on some event or condition and be published to the cloud independently of data collection. Data filtering should be available through basic analytics. Lastly, user setup should be minimalized by automating as much configuration as possible.

As industry continues to define IIoT, the concepts and realization of the optimal Embedded IIoT Solution will continue to evolve.

Appendix

Emerging IoT Standards

Protocol	Data Flow	Store and Forward	Performance	Adoption	Notes
AMQP	Pub/Sub over TCP/IP	✓	Seconds	Cloud	• Rich Connection and User Security • Packet Inspection for Routing (metadata)
MQTT	Pub/Sub over TCP/IP	✗	Seconds	Cloud	• Low overhead and embeddable • Lightweight user security
CoAP	Request/ Response	✗	Milliseconds	Edge	• Lightweight and connectionless • Optional Datagram Transport Layer Security (DTLS) for packet encryption
DDS	Open Pub/ Sub	✗	Microseconds	Edge	• Highly performant • Rich security

About Kepware, Inc.

Kepware Technologies is a private software development company headquartered in Portland, Maine. Kepware provides a portfolio of software solutions to help businesses connect diverse automation devices and software applications. From plant floor to wellsite to windfarm, Kepware serves a wide range of customers in a variety of vertical markets including Manufacturing, Oil & Gas, Building Automation, Power & Utilities, and more. Established in 1995 and now distributed in more than 100 countries, Kepware's software solutions help thousands of businesses improve operations and decision making. Learn more at www.kepware.com.

Connect With Us

Kepware wants to hear from you. Email Iot@kepware.com to share the challenges your operation faces in fulfilling the promise of the IoT.

Glossary

Asynchronous Message Queuing Protocol (AMQP)

A protocol primarily used to exchange data between gateway technologies in the cloud. It is a connection-oriented protocol over TCP/IP. It supports store and forward and embraces pub/sub data flows from the source to the destination.

Constrained Application Protocol (CoAP)

A request/response (or master/slave) protocol intended for use in simple devices for communication over the Internet. It is lightweight, and uses Hypertext Transfer Protocol (HTTP) over User Datagram Protocol (UDP).

Edge

The part of the network that bridges the gap between Information Technology and Operation Technology, where the rich resources available in the cloud are not directly available.

Industrial Internet of Things

The concept of connecting all hardware and software components within an industrial environment for complete visibility into operations at any time of day and from any location.

Information Technology (IT)

A business network (usually public and leveraging the Internet) consisting of hardware and software that allows for the storing, retrieving, and sending information for human consumption.

IPv4 Standard

The most widely-used version of the Internet Protocol (IP) that is used to connect devices to the Internet. It is able to uniquely identify 4 billion addressable items.

IPv6 Standard

The most recent version of the Internet Protocol (IP) that was designed to support the growth of the internet in terms of devices and data traffic. It will be able to identify 340 undecillionth (340 x 1036) addressable items.

Machine-to-Machine (M2M)

The ability for one machine to be able to connect to another and request its information.

Message Queuing Telemetry Transport (MQTT)

A protocol geared towards the embedded space that is used over TCP/IP and embraces pub/sub data flows from the source to the destination. Unlike AMQP, it is only capable of storing the last message for future transmission.

Operational Technology (OT)

An automation network (usually private) consisting of hardware and software that is responsible for monitoring and/or controlling industrial devices, processes, and events in the enterprise.

Predictive Maintenance

This technique includes monitoring operational and maintenance data to predict failures before they occur—providing the ability to schedule maintenance during planned downtime.

Process Optimization

This iterative process continually improves operational efficiency—reducing waste and saving costs.

Pull Model

A request/response model in which a Thing requests a piece of data over a communications medium and waits for a response containing the result.

Push Model

A pub/sub model in which a Thing pushes data over a communications medium and parties interested in the data subscribe to the stream.

Supervisory Control and Data Acquisition (SCADA)

This system provides remote visibility—enabling users to monitor and control remote data sources (or Things) over multiple sites and long distances.

Things

The hardware and software components that comprise an automation system.

Zettabytes

A unit of information equal to one sextillion (10^{21}) or 270 bytes.

Index

Kepware Technologies | 400 Congress Street | Portland, Maine 04101 | +1-888-KEPWARE | www.kepware.com