AUTOMATION 2023

VOLUME 3

IIoT & Industry 4.0

- The Modern, Smart Factory
- Taking the IIoT to Remote Sites
- Unlocking Digital Value
- Today's Human Machine Interfaces
- The Power of OEE Data
- MQTT for IIoT Challenges
- Distributed Data Lake Architecture

AUTOMATION.COM | ISA A subsidiary of the International Society of Automation

Connect The Dots With Ignition

The Unlimited Platform for Total System Integration

ALL YOUR
DATA

ALL YOUR
DEVICES

ALL YOUR
OPERATIONS

ALL YOUR
PEOPLE

ALL YOUR
LOCATIONS

ONE PLATFORM TO CONNECT ALL YOUR PROCESSES, PEOPLE & PROGRAMS

Learn more
about Ignition:
ia.io/platform

inductive
automation

Table of Contents

AUTOMATION 2023 VOLUME 3
IIOT & INDUSTRY 4.0

Introduction

AUTOMATION 2023 VOLUME 3

Leaving Room to Grow with IIoT & Industry 4.0

This May edition of AUTOMATION 2023 unfolds with a look at how modern, smart factories are built with unknowable future technologies in mind. A truly sustainable environment starts with early planning that allows for flexibility far down the road. With that in mind, we've brought you pieces from industry experts who share what they know about readers' very specific challenges: corralling all types of data and making it useful (and accurate) with a data mesh; HMIs continuing to provide operators with data and insights, even with the unrivaled connectivity we're seeing today; and much more. You'll also hear from the decision-makers who allow all these implementations to become a reality—it's not *all* about the technology. With this issue of AUTOMATION, we invite you to celebrate your wins along the way to Industry 4.0 and beyond as the innovation and integration of IIoT evolves.

If you have a moment, let us know what you think of this issue. We are always striving to bring our readers more of what they want and need to see in upcoming AUTOMATION ebooks.

Lynn DeRocco
Automation.com Managing Editor
lderocco@automation.com

About AUTOMATION 2023

The AUTOMATION 2023 ebook series covers Industry 4.0, smart manufacturing, IIoT, cybersecurity, connectivity, machine and process control and more for industrial automation, process control and instrumentation professionals. To subscribe to ebooks and newsletters, visit: www.automation.com/newslettersubscription.

AUTOMATION 2023 is published six times per year (January, March, May, July, September, and November) by Automation.com, a subsidiary of International Society of Automation (ISA). To advertise, visit: www.automation.com/en-us/advertise.

AUTOMATION.COM

groups/68581

automationdotcom

@automation_com

International Society of Automation
Setting the Standard for Automation™

company/internationalsocietyofautomation

InternationalSocietyOfAutomation

@ISA_Interchange

Renee Bassett, Chief Editor
rbassett@automation.com

Chris Nelson, Advertising Sales Rep
chris@automation.com

Richard T. Simpson, Advertising Sales Rep
rsimpson@automation.com

Gina DiFrancesco, Advertising Sales Rep
GDiFrancesco@automation.com

Holistic Planning for a More Modern Smart Factory

Built-in flexibility allows for changes over time as new technologies become available.

By Richard Slaugenhaupt and Lynn Njaa,
MAVERICK/Rockwell Automation

Today's ultra-competitive industrial climate has manufacturers looking for ways to become more agile and adapt quickly to rapidly changing markets and an increasingly diverse range of customer demands. Digital transformation is one approach that many enterprises are using to become more agile. To digitally transform their facilities, they are turning to smart technologies such as industrial internet of things (IIoT), digital twin/thread, artificial intelligence, machine learning, cloud computing, mobility, analytics, and so much more. For many, the ideal facility is a place where operations, IT, supply chain, and customers all coexist in an environment of collaboration and information sharing that breeds efficiency and reliability.

A good design helps to future-proof the automation solution to accommodate new features, functions, and technologies. These systems are likely to be around for decades, and the chance that changes and improvements will be needed and desired is pretty much a sure thing.

The road to a smarter, more modern facility isn't easy, however. For many manufacturers, a substantial gap exists between their current state and what is required to even begin adopting these technologies. Recognizing how smart technologies enable bottom-line value is a good first step toward bridging the gap, but recognition alone isn't sufficient to power a company's effort to competitively reach "elite" status. Many manufacturers struggle against an outdated control system infrastructure that limits their ability to adopt these important enabling technologies.

Before success can be realized, manufacturers must address these legacy restrictions and develop a capable foundation that supports new and emerging technologies. They should also take a hard look at their operations and see where smart technologies could fit into their existing infrastructure while continuing to support their business needs and process requirements. Ultimately, the road to a smart factory is a journey that requires innovation, an upfront plan, and a solid implementation strategy that is based on a holistic design approach with an eye toward modularity, agility, and ultimate sustainability.

A time to innovate

Tackling an automation hardware and software project of any magnitude is not for the faint of heart. Key stakeholders must decide whether to replace systems piecemeal or start over completely with all new systems. No matter the chosen path, a qualified team using a standards-based approach is critical to ensure consistency and meet company-wide functionality requirements.

As the heart of the production facility, the process control system is critical to business operations and, as such, is rarely altered. A business-sanctioned modernization project is the ideal opportunity to evaluate current operations and leverage smart technologies to not just replicate these critical systems—but innovate instead.

Replicating existing issues will reduce the flexibility of the new control system from the very outset and will continue to limit its potential over its lifespan. Manufacturers could end up back where they started—with the same functionality simply being implemented in different ways, creating maintenance challenges down the road.

This is not to say that manufacturers should outright dismiss the value of their existing legacy systems. It is still essential to preserve and leverage the years of intellectual property configured into existing systems. All their positive aspects should be captured and combined with the new technologies being implemented.

Incorporating the latest control functionality—developing effective human-machine interfaces (HMIs), alarm management, and compliant safety systems—will enhance profitability. The key to success is to identify what new technologies will provide the most benefit and return on investment (ROI) for a company's unique situation.

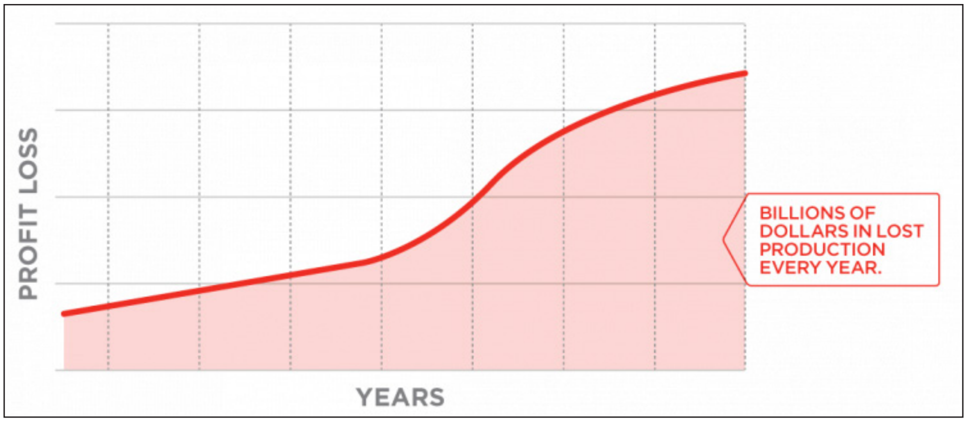

A multitude of factors, such as unaddressed legacy HMI issues, leads to poor ROI.

For instance, legacy HMIs are an area that has been consistently identified as a problem for operators. Unclear and ineffective graphics, along with an alarm system that lacks proper prioritization, improperly set alarm points, and ineffective annunciation lead to unplanned downtime and billions of dollars in lost production every year.

Facilities that deal with these real-world issues could benefit from a highly functioning alarm management system. A modern high-performance HMI can help operators respond more quickly to problems by clearly distinguishing between low-priority, high-priority, and critical alarms to mitigate abnormal situations before production is impacted.

These projects are inherently complex, and internal resources have limited bandwidth, so engaging an automation solutions provider with expertise in planning, design, and implementation can move a project forward with reduced risk and greater success. Depending on size, the project scope may need to be parsed into multiple projects, where each has its own schedule that includes:

- Current and subsequent front-end loading (FEL) process
- Design
- Engineering
- Testing
- Installation
- Startup and commissioning

A trusted third-party automation solutions provider can help evaluate the overall situation and create a well-defined project scope and framework that aligns with the company's business goals and objectives. It all starts with a solid plan.

Plan early

Front-end loading/front-end engineering (FEL/FEED): A wise and necessary first step for a modernization project is the FEL/FEED effort, which ensures that an execution plan and schedule are in place to keep the project(s) on time and on budget. Each phase of the FEL process (FEL1>FEL2>FEL3), from conceptual "go/no-go" for the project to capital planning estimates to refining the details that justify funding, helps reduce implementation risk for cost, schedule, and system performance. The FEL process also helps determine the best control system platform to meet the business needs and provides accurate, justifiable cost estimates to successfully navigate the capital expense (CapEx) funding process and secure stakeholder support.

In addition to technical aspects of an FEL, companies may benefit from an upfront CapEx evaluation and iterative implementation approach to balance improvement efforts with cash flow constraints. When strategic improvement goals are bound to a piecewise project execution method, however, it is important not to lose sight of desired long-term outcomes. Often, too much emphasis is placed on traditional objectives like simpler, cheaper, and faster evaluated on a per-project basis. This myopic approach can negatively impact the long-term goals of a broader modernization effort.

Project execution plan (PEP): Preparing and planning for startup and commissioning early in the process helps synchronize construction activities, which is extremely important to identify any potential equipment issues and minimize risk for a more successful project outcome. Prior to creating a plan, personnel should identify and define any potential risks; e.g., downtime, network traffic levels, data integrity, operator graphics and alarms, and determine how to mitigate them.

A detailed PEP for both the engineering and the installation of the new system includes reverse-engineering the functionality of the legacy system to ensure that none of the functionality is left behind and determining how the new system will be cutover, hot or outage. A PEP is also the time to gather information regarding piping and instrumentation diagrams (P&IDs), instrument specifications, single line diagrams, loop sheets, and other pertinent design documents. It should include well-defined documents for roles and responsibilities, a risk assessment, quality planning, testing planning, training plans, a resource-loaded schedule, and other appropriate items, all of which are important for effective execution.

Best practices, like holding regular project team meetings and following a proper communication plan, will help keep everyone on the same page and performing as planned. Document control will ensure that documentation is properly managed and updated so that the latest revisions are readily available and easily retrieved by the people using them. All of this will greatly improve efficiency when construction and commissioning activities begin.

Overall, proper upfront planning requires a holistic look at the entire facility's operations. All company-wide factors should be discussed before developing a migration plan while there is still the greatest flexibility to deal with them. It can't be emphasized enough—early effort spent on good planning and definition pays for itself many times over.

Smart data needs flexible networks

Networks are another area that requires attention early in the upfront planning process. Any attempt to modernize a company's operational proficiency should always start with a thorough evaluation of its network infrastructure.

The integrated technologies underpinning the smart factory run on data. Smart technologies are essential tools for capturing and analyzing data, and turning it into intelligence that enables data-driven decisions in real time. Recognizing that data is the lifeblood of the smart factory, and that data runs through networks, don't overestimate the importance of a capable and efficient network to operational success.

Many facilities run on networks that evolved, rather than resulted from thoughtful planning. Those networks often started small and grew as needed to support new operations or system functionality.

The unplanned sprawl that results from this sort of evolutionary growth produces networks that are sluggish and cumbersome to maintain.

Think of a network infrastructure as the central nervous system of the body, connecting all the parts together in a way that enables them to function as a unit. Without it, facilities just have a collection of disjointed parts. Like the nervous system, a proper networking infrastructure operates efficiently and homogeneously, even when different mechanisms are employed. Like the body uses both electrical and chemical messages for passing information, a network will utilize multiple protocols to achieve the best fit of features to function. Replacing outdated or hard-to-support protocols with mainstream ones will minimize incompatibility problems while shoring up the enterprise network.

Many facilities run on networks that evolved, rather than resulted from thoughtful planning. The unplanned sprawl that results from this produces networks that are sluggish and cumbersome to maintain.

Design and develop for agility

A good holistic design approach is foundational for addressing the long-term evolving needs of a smart factory. Success results from production processes and systems that can adapt quickly to change, not ones hindered by cumbersome designs and slow change management procedures.

Legacy control systems were designed as stand-alone solutions, without standardization of method or function. When considered all together, these systems often have holes and/or overlap in the functionality required for them to interoperate efficiently. To provide a proper foundation for the smart factory, these systems would need to have been constructed in a cohesive manner from a master plan or template, which is rarely the case.

Depending on a facility's infrastructure, the approach to the redesign effort varies, but the best design will blend appropriate quantities of tried-and-true methods with value-adding modern technologies. Modular designs support seamless upgrades using standards-based products and software. Object-oriented design, for example, is one technology that can greatly improve the speed and agility of project implementation while establishing a consistent, flexible, and capable platform for handling any of the smart factory's functional needs.

Implementation and execution

When it comes time for the implementation phase of an automation project, qualified personnel will need to execute rigorous testing and commissioning procedures to ensure safe and efficient startup. This is where a qualified team will use the gathered documentation from the PEP to create dynamic test plans and document procedures, such as loop checks, control system hot cutovers, and more.

These detailed plans serve as the execution blueprint for each pre-commissioning assessment and are used to verify that each device performs both individually and systematically as intended prior to startup.

Shortcuts here will inevitably lead to operational problems, costing many times over any savings realized from reduced commissioning time. Take the time to follow thorough procedures created with input from subject matter experts who have a track record of successful project execution.

The path forward: sustain and improve

Much emphasis is placed on a solid plan and a holistic design, as these important elements greatly impact the success of implementing and sustaining a process control system. The details of the execution plan ultimately determine the timeless viability of a new automation solution based on that initial design. If planned for from the start of the project, continuous improvement initiatives can help manufacturers optimize processes and increase operational safety and efficiency, giving them an ongoing competitive advantage.

It is important to establish standards and good automation engineering best practices to ensure that the new automation system is designed and configured in a way that meets your facility's requirements while being maintainable over the long term. A good design helps to future-proof the automation solution to accommodate desired new features, functions, and technologies. These systems are likely to be around for decades, and the chance that changes and improvements will be needed and desired during a system's lifetime is pretty much a sure thing.

To this end, as manufacturers plan, develop, and build their smart factories, they must also consider new approaches, such as setting up secure remote infrastructures and establishing remote application support capabilities, to continuously optimize, maintain, and sustain them.

Where remote access capabilities are required, legacy systems present a challenge as hardware and software support, security updates, and software fixes for older releases become obsolete. Manufacturing personnel will need to safeguard their IP with critical updates, reboots, and patches for system security and reliability. With this in mind, manufacturers should develop remote management strategies whereby they leverage smart technologies, mobile devices, remote access connectivity, and communications tools to ensure that critical infrastructures remain up and running efficiently.

Reducing the time in-house personnel must perform critical automation tasks can realize a cost savings. Consider outsourcing remote management and preventative maintenance services to help ease the workload.

Reducing the time in-house personnel must perform or be trained on critical automation tasks can realize a cost savings. Consider outsourcing remote management and preventative maintenance services to help ease the workload for personnel having to perform a variety of critical tasks, including incident management, system backup and recovery, data analysis and retrieval, software updates and patches, real-time monitoring of software and systems, machine learning, predictive analytics, IT/OT hardware refresh, online automation, and system programming edits.

Regardless of a facility's size, the road to a smart factory is navigable if there's a clear path outlined at the beginning. A modernization project's successful outcome depends heavily on a disciplined approach and a holistic plan from initial design concept through implementation and startup. Long-term goals to sustain and continuously improve the new smart factory further ensure flexible, agile systems for ongoing competitive success.

Images courtesy of Rockwell Automation

ABOUT THE AUTHORS

Richard Slaugenhaupt is a consultant for both MAVERICK Technologies, a leading automation solutions provider, and Rockwell Automation. He has been working in manufacturing automation for 40 years, serving both small and large customers in many different industries.

Lynn Njaa is a senior proposal consultant, lifecycle services for both MAVERICK Technologies, a leading automation solutions provider, and Rockwell Automation. Lynn helps guide the strategic vision on automation projects and execution management. She brings diverse automation consulting experience in systems engineering, proposal and estimate engineering, business development, project management, and commercial development.

PROVEN

40

YEAR

OPERATING

LIFE*

Remote wireless devices connected to the Industrial Internet of Things (IIoT) run on Tadiran bobbin-type $LiSOCl_2$ batteries.

Our batteries offer a winning combination: a patented hybrid layer capacitor (HLC) that delivers the high pulses required for two-way wireless communications; the widest temperature range of all; and the lowest self-discharge rate (0.7% per year), enabling our cells to last up to 4 times longer than the competition.

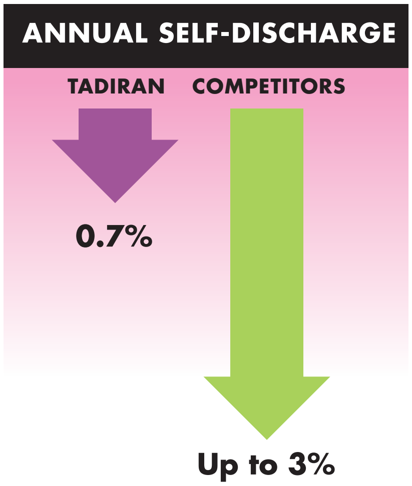

Looking to have your remote wireless device complete a 40-year marathon? Then team up with Tadiran batteries that last a lifetime.

TADIRAN BATTERIES

Tadiran Batteries
2001 Marcus Ave.
Suite 125E
Lake Success,
NY 11042
1-800-537-1368
516-621-4980

www.tadiranbat.com

* Tadiran $LiSOCL_2$ batteries feature the lowest annual self-discharge rate of any competitive battery, less than 1% per year, enabling these batteries to operate over 40 years depending on device operating usage. However, this is not an expressed or implied warranty, as each application differs in terms of annual energy consumption and/or operating environment.

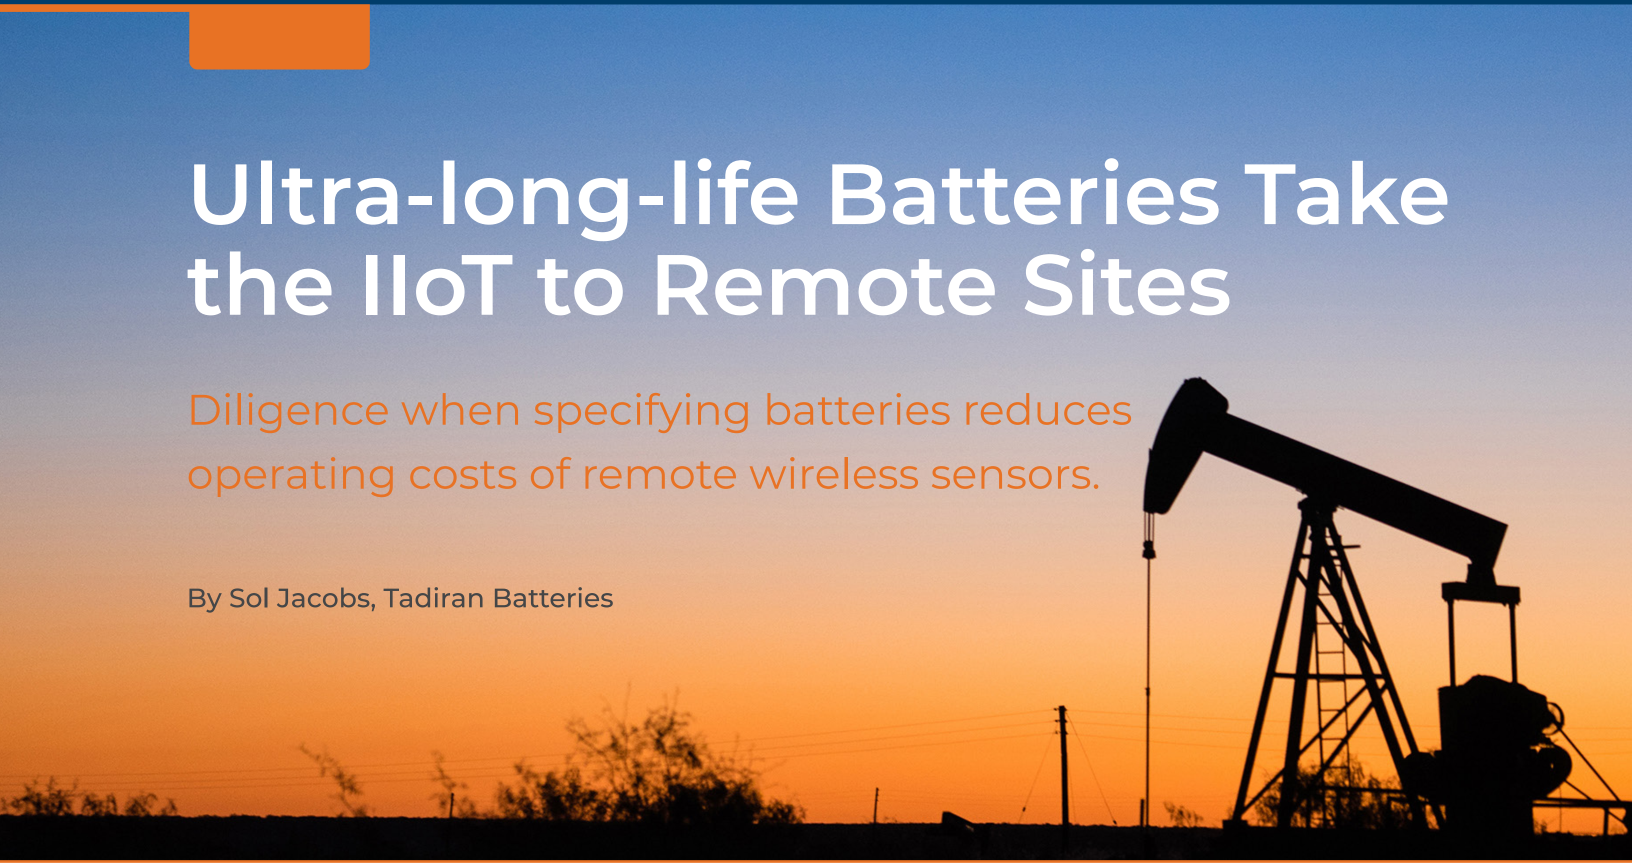

Ultra-long-life Batteries Take the IIoT to Remote Sites

Diligence when specifying batteries reduces operating costs of remote wireless sensors.

By Sol Jacobs, Tadiran Batteries

The Industrial Internet of Things (IIoT) is expanding to increasingly remote locations, with lithium battery-powered remote wireless devices bringing digital connectivity to virtually all industrial applications, including SCADA, process control, industrial robotics, asset tracking, safety systems, environmental monitoring, M2M, AI, and wireless mesh networks, to name a few.

Industrial-grade lithium batteries enable remote data to be applied more intelligently to improve operational efficiency, enhance quality control, track assets, promote greater environmental sustainability, optimize supply chains, enhance in-field predictive maintenance programs, and more. Use of batteries also eliminates the expense and time-consuming task of having to hard-wire the devices.

With numerous battery chemistries to choose from, the process of identifying the ideal power involves various criteria, including:

- Determining the energy demand
- Identifying the ideal battery chemistry

- Understanding the importance low battery self-discharge
- Adapting the solution for high pulse requirements
- Comparing batteries with similar chemistries

Determining the energy demand

A remote wireless device is only as reliable as its battery. In order to maximize operating life, design engineers must consider numerous factors such as the amount of energy consumed during active mode (including the size, duration, and frequency of pulses); the amount of energy consumed while the device is in standby mode (the base current); the length of storage (as normal self-discharge during storage diminishes capacity); the impact of thermal environments (including

Ayyeka AI-enabled smart sensors are used to monitor and maintain hard infrastructure used in solid waste and wastewater management, public utilities, transportation, energy exploration and distribution, smart cities, environmental monitoring, and more. Powered by Tadiran bobbin-type $LiSOCl_2$ batteries, these remote wireless devices provide two-way wireless communications to maximize operational efficiency, detect unusual events, support predictive maintenance programs, and counter cyber security threats.

storage and in-field operation); equipment cut-off voltage (as battery capacity becomes exhausted, or in extreme temperatures, voltage can drop to a point too low for the sensor to operate). Most critically, the design engineer must consider the annual self-discharge rate of the battery, which often exceeds the amount of energy consumed while operating the device.

If the application is easily accessible for battery replacement and the operating environment is relatively moderate, the solution could be as simple as utilizing an inexpensive consumer-grade alkaline or lithium battery. Conversely, if the application involves a long-term deployment in a hard-to-access location or extreme environment where battery replacement is prohibitively expensive or impossible, an industrial-grade lithium battery is generally required.

With numerous battery chemistries to choose from, the process of identifying the ideal power involves various criteria.

To conserve energy and extend operating life, low-power IIoT devices operate in a standby state, drawing micro-amps of average current. These devices may also require periodic high pulses in the multi-amp range to power bi-directional wireless communications.

Low-power IIoT devices are predominantly powered by bobbin-type lithium thionyl chloride ($LiSOCl_2$) batteries that feature very high capacity, high energy density, an extended temperature range, and an exceptionally low annual self-discharge rate. Certain niche applications draw higher amounts of average current measurable in milli-amps with pulses in the multi-amp range, which is enough average energy to prematurely exhaust a primary (non-rechargeable) battery. These specialty applications may be better suited for an energy harvesting device in combination with an industrial-grade Lithium-ion (Li-ion) battery to store the harvested energy.

Primary Cell	$LiSOCL_2$ Bobbin-type with Hybrid Layer Capacitor	$LiSOCL_2$ Bobbin-type	Li Metal Oxide Modified for high capacity	Li Metal Oxide Modified for high power	$LiFeS_2$ Lithium Iron Disulfate (AA-size)	$LiMnO_2$ Lithium Manganese Oxide
Energy Density (Wh/Kg)	700	730	370	185	335	330
Power	Very High	Low	Very High	Very High	High	Moderate
Voltage	3.6 to 3.9 V	3.6 V	4.1 V	4.1 V	1.5 V	3.0 V
Pulse Amplitude	Excellent	Small	High	Very High	Moderate	Moderate
Passivation	None	High	Very Low	None	Fair	Moderate
Performance at Elevated Temp.	Excellent	Fair	Excellent	Excellent	Moderate	Fair
Performance at Low Temp.	Excellent	Fair	Moderate	Excellent	Moderate	Poor
Operating Life	Excellent	Excellent	Excellent	Excellent	Moderate	Fair
Self-Discharge Rate	Very Low	Very Low	Very Low	Very Low	Moderate	High
Operating Temp.	-55°C to 85°C, can be extended to 105°C for a short time	-80°C to 125°C	-45°C to 85°C	-45°C to 85°C	-20°C to 60°C	0°C to 60°C

Table 1. Lithium battery types.

Identifying the ideal battery chemistry

Numerous primary (non-rechargeable) lithium battery chemistries are available (Table 1). At one end of the spectrum are inexpensive alkaline batteries that deliver high continuous energy but suffer from a very high self-discharge rate (which limits battery life), low capacity and energy density (which adds size and bulk), and an inability to operate in extreme temperatures due to the use of water-based constituents. At the opposite end of the spectrum are industrial grade lithium chemistries.

As the lightest non-gaseous metal, lithium features an intrinsic negative potential that exceeds all other metals, delivering the highest specific energy (energy per unit weight), highest energy density (energy per unit volume), and higher voltage (OCV) ranging from 2.7 to

3.6V. Lithium battery chemistries are also non-aqueous, and therefore less likely to freeze in extremely cold temperatures.

Bobbin-type lithium thionyl chloride ($LiSOCl_2$) batteries are overwhelmingly preferred for long-term deployments since they deliver the highest capacity and energy density, endure the most extreme temperatures (-80°C to +125°C), and feature an annual self-discharge rate as low as 0.7% per year for certain cells, thereby creating the potential for 40-year battery life. Bobbin-type $LiSOCl_2$ batteries offer the following benefits:

- **Higher reliability**—This is ideal for remote locations where battery replacement is difficult or impossible, and highly reliable connectivity is required.
- **Long operating life**—Since the battery's self-discharge rate often exceeds actual energy usage, high initial capacity and a low self-discharge rate are highly beneficial.
- **The widest temperature range**—Bobbin-type $LiSOCl_2$ cells can be modified to operate reliably in extreme temperatures (-80°C to 125°C).
- **Smaller size**—Higher energy density may permit the use of smaller batteries.
- **Higher voltage**—This could allow for the use of fewer cells.
- **Lower lifetime cost**—Can be a major consideration since the manpower and logistical expenses required to replace a battery will far exceed its cost.

Cattlewatch AI-enabled electronics collars allow ranchers to remotely track their cattle herds by providing behavioral information and alerts using an ultra-low-power LoRaWAN network. Select members of the herd are equipped with solar-powered communicators that form a wireless mech network involving the entire herd. Tadiran TLI Series rechargeable li-ion batteries create a lightweight solution that can withstand extreme temperatures, offers up to 20-year operating life and 5,000 full recharge cycles, and generates the high pulses required to power remote wireless communications.

The importance of low battery self-discharge

All batteries experience some amount of self-discharge as chemical reactions draw small amounts of current even while the cell is unused or disconnected.

Self-discharge can be minimized by controlling the passivation effect, whereby a thin film of lithium chloride (LiCl) forms on the surface of the lithium anode to separate it from the electrode to reduce the chemical reactions that cause self-discharge. Whenever a load is placed on the cell, the battery experiences initial high resistance and a temporary drop in voltage until the discharge reaction begins to dissipate the passivation layer: a process that keeps repeating every time a load is applied.

The passivation effect can vary based on the cell's current discharge capacity, the length of storage, storage temperature, discharge temperature, and prior discharge conditions, as partially discharging a cell and then removing the load increases the level of passivation over time. While harnessing the passivation effect is essential to reducing self-discharge, too much it can be problematic if it overly restricts energy flow.

Bobbin-type $LiSOCl_2$ cells vary significantly in terms of their ability to harness the passivation effect. For example, top quality bobbin-type $LiSOCl_2$ batteries can feature a self-discharge rate as low as 0.7% per year, thus retaining nearly 70% of their original capacity after 40 years. Conversely, lower quality $LiSOCl_2$ cells can have a self-discharge rate as high as 3% per year, exhausting nearly 30% of their available capacity every 10 years, which greatly reduces their operating life.

Adapting the solution for high pulse requirements

Low-power remote wireless devices increasingly require periodic pulses up to 15 A to support two-way wireless communications. Standard bobbin-type $LiSOCl_2$ cells are unable to deliver these high pulses due to their low-rate design. This hurdle can be easily overcome with the addition of a patented hybrid layer capacitor (HLC). This hybrid solution utilizes the standard bobbin-type $LiSOCl_2$ cell to deliver nominal background current during standby mode while the HLC delivers high pulses to support data transmission. As an added bonus, the HLC experiences a unique end-of-life voltage plateau that can be interpreted to generate low battery status alerts.

Supercapacitors perform a similar function with consumer electronics but are poorly suited for industrial applications due to serious limitations, including: short-duration power; linear discharge qualities that do not allow for the use of all available energy; low capacity; low energy density; and very high self-discharge rates up to 60% per year. Supercapacitors linked in series also require the use of expensive cell-balancing circuits that add bulk and drain additional current to further shorten their operating life.

Comparing seemingly similar batteries

To maximize return on investment (ROI), the ideal battery-powered solution should last for the entire lifetime of the device to reduce or eliminate the need for costly battery change-outs. However, it can take years to differentiate a higher quality battery from a poorer quality cell since the initial capacity losses are not easily measurable. In addition, the theoretical models and algorithms used to calculate battery life expectancy can be highly unreliable since they tend to underestimate the passivation effect as well as long-term exposure to extreme temperatures.

Careful due diligence is required when specifying an ultra-long-life battery. All potential suppliers should be required to provide fully documented and verifiable test results along with in-field performance data involving similar devices operating under similar loads and environmental conditions. Going the extra mile to carefully compare batteries could pay important dividends by increasing product longevity and lowering the total cost of ownership.

Images courtesy of Tadiran Batteries

ABOUT THE AUTHOR

Sol Jacobs is the vice president and general manager of Tadiran Batteries. He has more than 30 years of experience in powering remote devices. His educational background includes a BS in engineering and an MBA.

Unlocking Value in the Digital Era

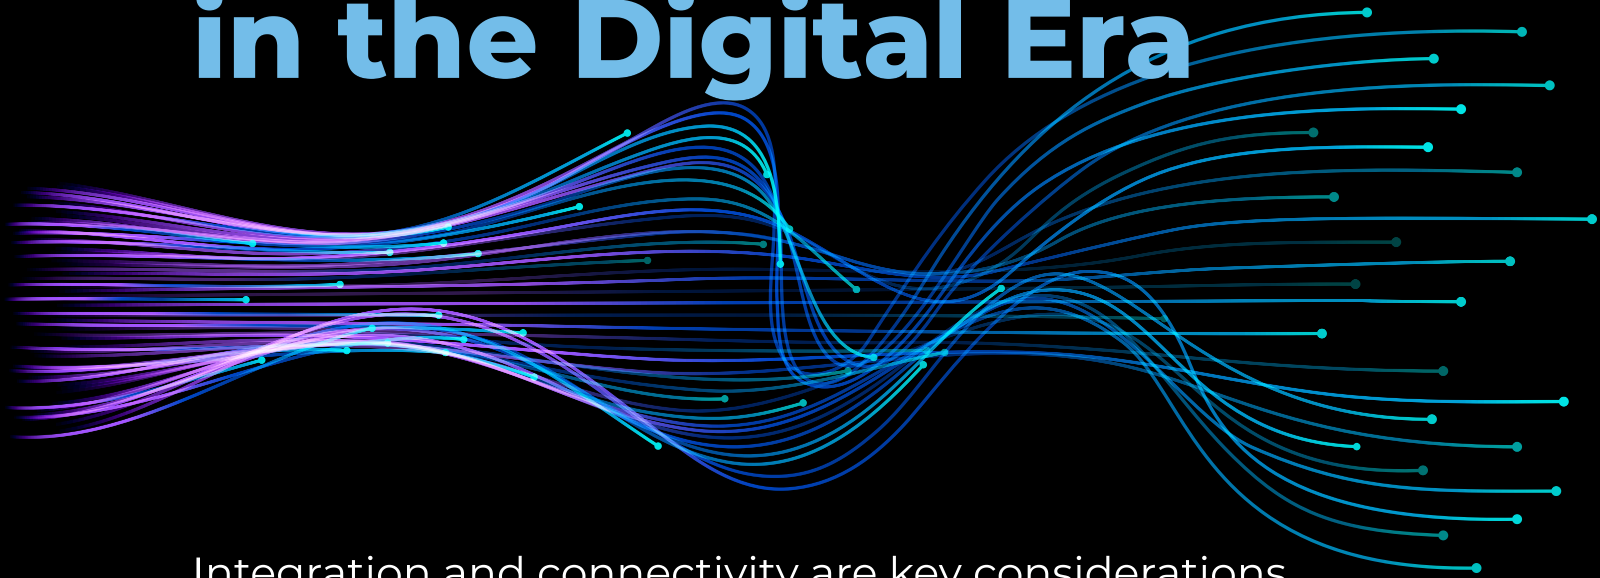

Integration and connectivity are key considerations in successfully employing Industry 4.0 solutions.

By Dawid Sadie,
Software Toolbox Inc.

Industry 4.0, also known as the Fourth Industrial Revolution, refers to the integration of advanced technologies such as the Internet of Things (IoT), artificial intelligence (AI), and big data into industrial processes. A key enabler of Industry 4.0 is the Industrial Internet of Things (IIoT), which describes the collection and analysis of vast amounts of data from sensors and devices on the factory floor. This data can then, in the correct format, be utilized to optimize processes, reduce downtime, and improve product quality.

In our experience, supporting users and integrators, device connectivity protocols, IT/OT orchestration, transactional system data integration, data cleansing, normalization, and contextualization are key supporting factors that must be considered.

Device connectivity

The scale and wide-reaching scope of Industry 4.0 industrial operations require efficient, performant, and inclusive communication technologies. The open protocols MQTT Sparkplug B and Open Platform Communication Unified Architecture (OPC UA) are leaders in addressing standardization and interoperability in IIoT. Our intention is not to determine the superior communication solution but rather to discuss their characteristics and aspects worth considering.

MQTT is a lightweight communications transport protocol that is suitable for limited bandwidth networks and applications with multiple clients and devices sharing data in a many-to-many arrangement. It enables clients to publish and subscribe to data in cloud- or premise-hosted brokers that manage the data.

MQTT alone does not define the organization of data in packets, known as the payload, leading to risks of interoperability issues and vendor lock-in due to vendor-specific payload formats. MQTT Sparkplug B extends basic MQTT with a standardized payload format for users, integrators, and suppliers to use to define models for interchanging data, albeit not as well defined as OPC UA information models.

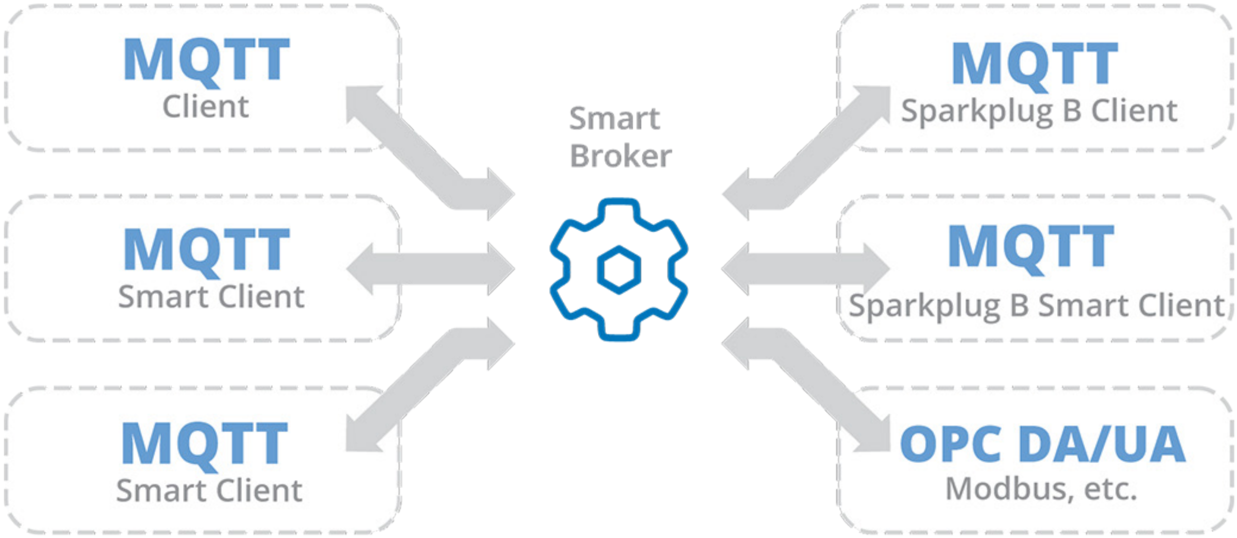

Smart brokers handle multiple payload formats over the same connection.

Users of MQTT need to consider Smart MQTT clients and brokers that go beyond just moving data around. A smart client or broker will handle multiple payload formats over the same connection, support emerging standards such as the IETF draft JSON schema in addition to Sparkplug B, be able to automatically extract MQTT topic data into tags for consumption via standards like OPC, and handle propagation of data connection quality status.

Strategies to connect and integrate standard PLC and control protocols, legacy devices, and nonstandard protocol devices must be addressed early, either through device replacement or integration software.

If control decisions rely on MQTT data, it is imperative for the smart client or broker to address the management of missed and out-of-sequence messages, guarantee message order preservation, and address the handling of failed writes. If network downtime and lost data is a concern, the smart broker or client must support store and forward. For robust security, the smart client and broker must enable designation of read-only data.

The OPC UA standards are another means of providing a standardized framework for data exchange and communication between diverse industrial systems, devices, and applications. An evolution of the OPC Classic standards, the OPC UA standards define a secure integrated means of exchanging a wide range of industrial data, along with standardized information models with well-defined namespaces for interchange of data, often in specific vertical industries.

OPC UA information models are published and available from the OPC Foundation with XML definition files to rapidly empower client and server applications to share the industry-specific data in the model. Sophisticated users and integrators can define and publish their own information models to exchange plant data within their businesses, or with supply chain or other partners. This data encompasses not only the raw data but also historical and event data and metadata, including

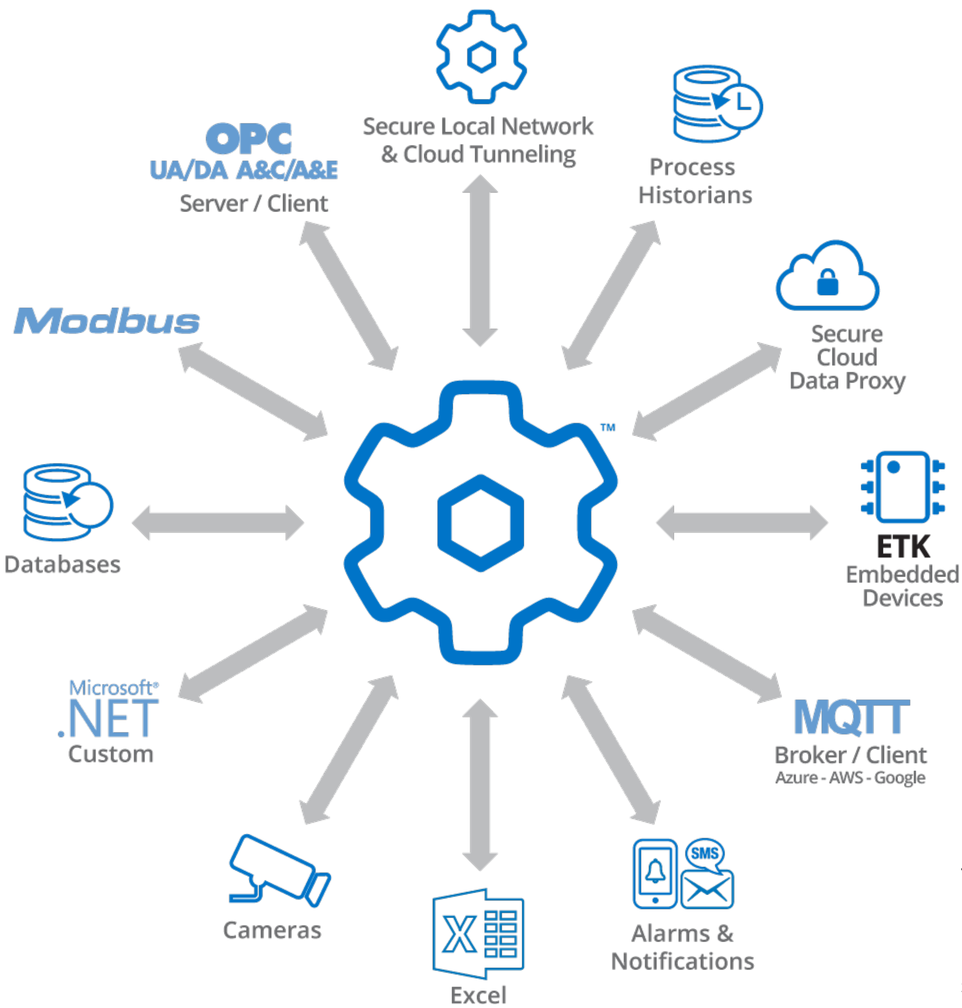

Middleware platform for bridging between proprietary device protocols and open standard protocols.

details about data sources, data quality, and interrelationships between data points.

If an OPC UA information model is available for your industry, you may consider using it. If your HMI, SCADA, historian, or MES does not yet support OPC UA information models, visual integration middleware solutions are available to transform the information model data into OPC UA data access or OPC Classic interfaces for integration into those systems.

For applications requiring many-to-many communications, OPC UA Publish/Subscribe (Pub/Sub) offers an efficient transport for data of all types, raw or organized, in OPC UA information models. OPC UA Pub/Sub is an extension of the OPC UA protocol. Unlike basic OPC UA, which uses a client/server model, it uses a secure, multicast-based model for simultaneous distribution of data to multiple subscribers. It can also be used over an MQTT transport, enabling businesses to benefit from the best of both worlds, leveraging the strong standardization of OPC UA information models with OPC Pub/Sub for data distribution, and the flexibility and simplicity of MQTT.

In any implementation, there will be systems and devices to connect that do not natively implement MQTT/OPC UA in any form. Strategies to connect and integrate standard PLC and control protocols, legacy devices, and nonstandard protocol devices must be addressed early, either through device replacement or integration software. The MQTT and OPC standards have empowered a market supply of off-the-shelf software with visual configuration interfaces to connect just about anything with a serial or ethernet connection and a documented, published protocol to the OPC and MQTT standards.

There is an active industry trend to create and enable a Unified Name Space (UNS) to empower efficient information interchange in support of real-time decision-making. The use of a UNS helps to break down information silos, allowing businesses to collect and analyze data from a wide range of sources, providing a more complete picture of operations. The UNS does not necessarily live in one place but rather in a distributed environment involving all applications.

The tools used to gather information that goes into the model should empower the standardization efforts. OPC UA, UA information models, UA Pub/Sub, MQTT, and MQTT Sparkplug B are technologies that meet those requirements, enabling and empowering technologies to get the data into the UNS. Middleware supported by suppliers with deep expertise in the field who are willing to have open conversations is critical to supporting UNS implementations as the bridge between proprietary device protocols and open standard protocols.

Ultimately, the decisions on what standards and tools are best for an application should stem from considerations about the nature of devices, the protocols they support, the protocols other software and business systems in the application support, and the availability of well-supported, off-the-shelf protocol conversion middleware. Other considerations include robust security, scale, performance, flexibility to accommodate future growth, and the existence of ecosystems of devices, software, and suppliers that are committed to supporting them.

Orchestrating OT and IT

The convergence of operational technology (OT) and information technology (IT) is also vital for dismantling data silos, facilitating uninterrupted data exchange, empowering real-time analysis, supporting predictive maintenance and quality control, and more. The modern capacity to automate processes reduces the risk of human error and lets businesses concentrate on more value-added activities. A flexible, event-driven middleware and visual workflow tool can play a pivotal role in achieving these goals.

Modern implementers expect integration without writing custom code and a visual working environment. They want to construct basic and complex workflows to exchange data and automate processes that are unique to their operations, test them, and then deploy them at scale. Since Industry 4.0 applications often involve massive amounts of data sources, the solution should scale using templates and mass imports. Users should expect a tool to support a wide array of protocols for OT, cloud, and IT, including OPC UA, OPC Classic, MQTT, REST, ERP interfaces, and databases

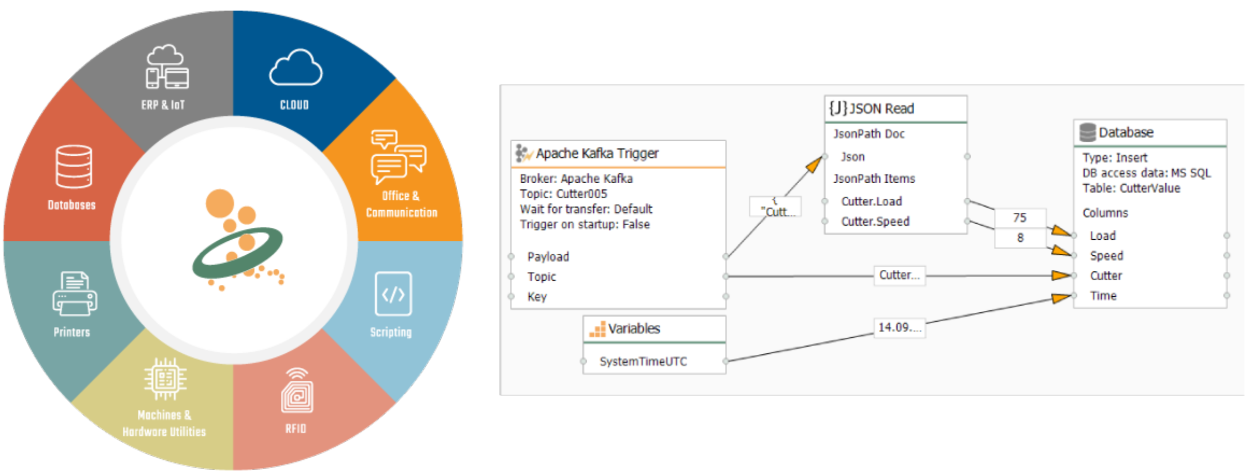

A flexible, event-driven middleware and visual workflow tool maximizes the possibilities for automation.

to ensure seamless data flow throughout the entire industrial ecosystem. By leveraging open standards-based tools such as OPC servers, they should also be able to reach IT systems using protocols such as SNMP.

When considering the choice and implementation of a visual workflow tool, businesses should reflect on the following aspects:

- **Current integration challenges:** Evaluate the IT and OT landscape to identify pain points that can deliver rapid time-to-value if eliminated, such as data silos, inefficiencies, or lack of real-time data access.
- **Find rapid time-to-value:** By starting with quick successes that bring meaningful business results, support and funding for the more complex longer-term gains should be easily obtained. A visual workflow tool will empower rapid prototyping and success, and templates and imports should be available for scaling up to the larger implementations funded by the early wins.
- **Scalability and adaptability:** A suitable tool should be scalable and adaptable in performance, deployment options, and configuration, allowing businesses to accommodate future growth, emerging technologies, and evolving industry standards.
- **Security and reliability:** The tool should ensure secure data transmission and storage and provide reliable performance in diverse industrial environments.
- **Ease of use and maintainability:** A user-friendly interface, troubleshooting tools, and robust support for various protocols and systems can reduce the learning curve and long-term maintenance efforts.

Once the devices are connected, and IT and OT technologies are orchestrated, there's still a key factor of bridging transactional operational and business data, cleansing, and contextualizing.

Bridging the data divide—transforming to value

Transforming data into information, also called the "data divide," adds value by providing actionable insights and enabling informed decision-making, driving operational efficiencies, and supporting continuous improvement initiatives. We discussed OT data access and OT to IT orchestration, but the integration of operational and business or transactional databases and systems also needs to be taken into account.

Bridging this so-called data divide requires a comprehensive solution that goes beyond the typical connectivity above. Essential functionality for such a solution includes cleansing, normalizing, and contextualizing, and providing data to various consumers. This work is more efficiently done closer to the data source, rather than sending potentially invalid or unaggregated data to the cloud for analytics, machine learning, and other value additions.

When evaluating such data normalization, data cleansing, information delivery, and bridging solutions, the following items must be considered:

- **Breaking down data silos:** Data silos occur when information is confined within separate systems, making it difficult to access and analyze. An ideal solution will integrate disparate systems,

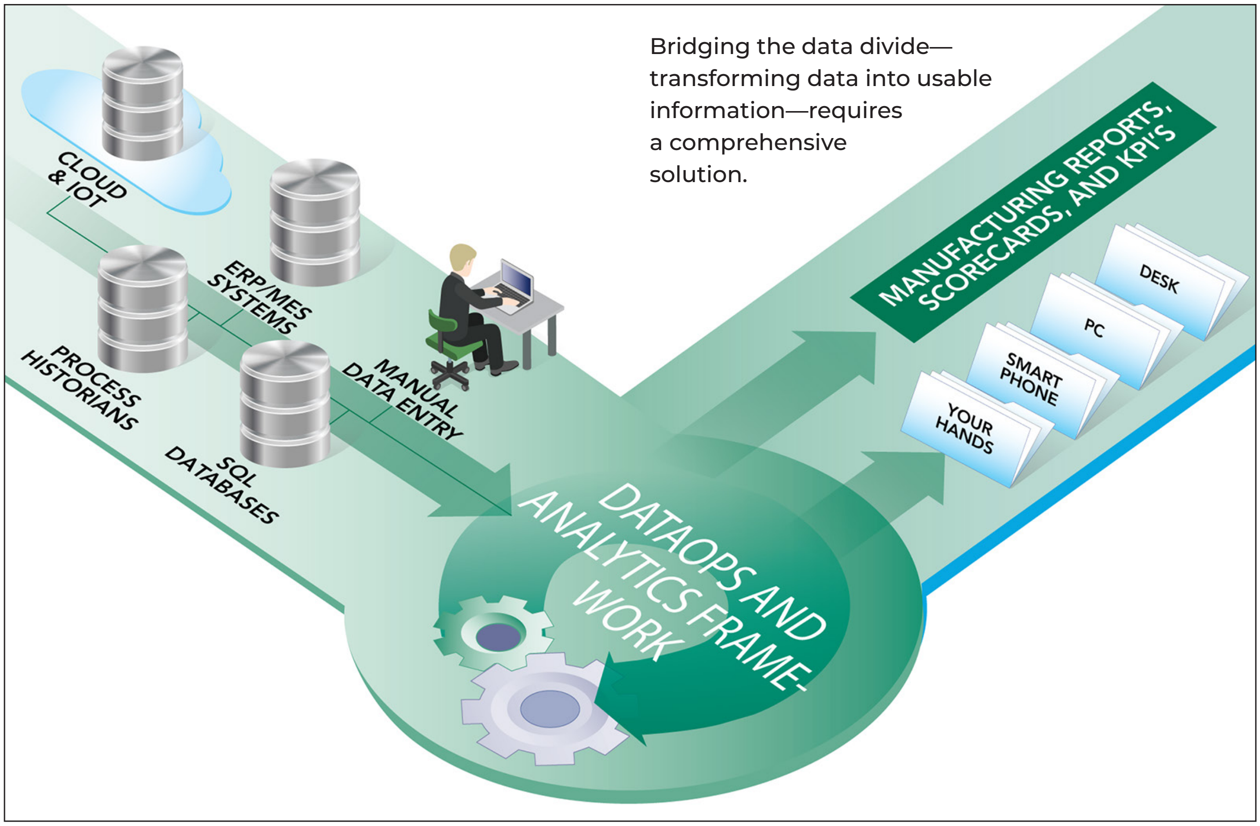

Bridging the data divide—transforming data into usable information—requires a comprehensive solution.

centralizing data from multiple OT, IT, and business sources into a unified platform. By providing a holistic view of the plant floor, the solution should enable organizations to optimize operations, identify inefficiencies, and make more informed decisions.

- **Streamlining decision-making and KPI delivery:** Advanced analytics and reporting tools are crucial for organizations to gain valuable insights from their data. Empowered operators and managers should not have to wait on developers to create the reports they need to answer their questions. Real-time self-service dashboards that provide key performance indicators (KPIs) and other critical metrics should be a part of the solution, empowering relevant team members to monitor performance, identify trends, and act.

- **Reducing manual data collection:** Manual data collection is time-consuming and prone to errors. The solution should automate data collection, reducing the need for manual intervention and ensuring accurate, consistent information. This empowers plant floor personnel to focus on more strategic tasks, increasing overall productivity. In the event where manual data collection is still required, the solution should provide an intuitive, browser-based, spreadsheet-like user interface to capture such data, complete with an audit trail and necessary security to protect data integrity.

- **Scalability and flexibility:** The solution should be designed to adapt to the unique needs of each organization. Ideally, the solution's modular architecture should allow for seamless integration with existing OT, IT, and business systems, and an ability to scale as business requirements evolve. The solution should support various communication protocols and interfaces, making it compatible with a wide range of devices and applications, including OT- and IT-related protocols, and transactional enterprise databases.

- **The value of data analytics:** Data analytics plays a crucial role in transforming raw data into actionable insights. An ideal solution will collect and centralize data near the source but also cleanse and contextualize it, ensuring that the information is accurate, relevant, and ready for analysis locally or in the cloud. Then, the solution can serve as a single source of truth for various data consumers, such as managers, engineers, executives, and other applications, that need to make data-driven decisions that improve the business.

Success in Industry 4.0 initiatives requires consideration of many factors. Device connectivity, OT to IT orchestration, bridging data divides to business/transactional systems along with cleansing and contextualization before performing advanced analytics are key enablers that must not be overlooked.

Software Toolbox stands as an experienced partner for clients and integrators seeking to address these key enablers. Established in 1996, the company offers an extensive array of open, standards-based tools that function collectively as solutions or as value-enhancing supplements to enterprise vendor applications. Our versatile offerings, including communication drivers, OPC server and client components and toolkits, and data visualization tools, allow for seamless integration of disparate systems, enabling clients to unlock the true value of their data. Software Toolbox's commitment to providing outstanding customer support and training ensures that clients have the necessary resources to implement, maintain, and expand their Industry 4.0 capabilities. Ultimately, Software Toolbox's expertise and dedication to client success serve as the foundation for the digital transformation that fuels the future of industrial automation.

Images courtesy of Software Toolbox Inc.

ABOUT THE AUTHOR

Dawid Sadie is a business developer at Software Toolbox Inc. with a 22-year background in automation technology. His experience ranges from working as an integrator to leading technical sales teams and heading a process automation partnership. Dawid offers his global expertise in IoT/IIoT, smart manufacturing, and advanced solutions. Now based in the US, he is dedicated to empowering client success through his extensive knowledge and hands-on experience.

Giving Modern **Human Machine Interfaces** Their Due

In a smart connected world, plants rely on HMIs for detailed data *and* industrial control insights.

By Charlie Norz, WAGO

Does the term Human Machine Interface (HMI) make you think of a graphic screen used by plant floor operators to help run a piece of equipment? Are you smiling because you know that now there is one screen that has replaced dozens of industrial push buttons, pilot lights, and 16-position selector switches marked with unrecognizable abbreviations? It's true, even though many people still have old-fashioned mental pictures of HMIs.

An HMI today is a screen that enables right-brained engineers the flexibility to create functional art that beautifully depicts a machine's operation. The operator can switch screens to command machine operation, supervise machine performance in widgets per hour, or view machine health via maintenance screens. The HMI is the human interface with the machine, or, more correctly stated, with the PLC that it is connected to. Plant floor information from the HMI and PLC is sent to a SCADA system that plant supervisors can then use to generate summary reports at the end of the week.

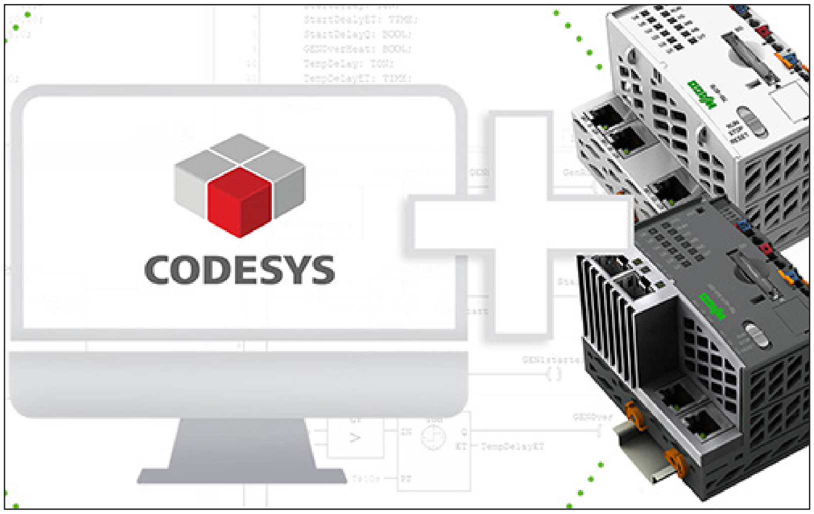

The CODESYS-based runtime environment and real-time-capable Linux operating system of WAGO's PFC200 Controller.

HMIs have successfully been employed for decades as key instruments used by plant floor operators to run, supervise, and maintain skidded equipment. They have saved manufacturers money by replacing dozens of 30mm push buttons, pilot lights, and selector switches that had needed to be wired to the PLC's I/O modules. HMIs have become an exciting way to operate machines on the plant floor, tank farm, engine room, packaging area, and many other applications.

Times are changing. In the smart connected world, manufacturers empower workers with knowledge to make decisions based on business-related data and its analysis. In this competitive era, the plant floor operators of yesterday are now the technical users of today. These individuals need to have status information about equipment to the east and west of their production line. They need the tools to predict what will be the raw material status at mid-shift and end of shift. They need to know the overall equipment effectiveness compared to the machine's designed specifications, compared to last night's shift, or twin plants in Europe and Asia, all in real time. Employers have found that providing business information to their staff enables them to make better decisions at all job functions and thus increases optimal manufacturing performance.

Companies that employ Industrial Internet of Things (IIoT) tools have been enjoying the many business benefits that come along with them. The data collected via the cloud allows the headquarter staff to view data about any plant, any production line, and any machine in the enterprise. However, if the data is static and remains at HQ only to be reported out at the end of the week, the opportunity to make dynamically integrated corrections or enhancements could be lost. Plant floor operators require live, current data in their hands to make real-time decisions. Companies with IIoT systems quickly realize that the benefit is not the data; it's that their staff has the tools to make well-informed decisions.

Controls engineers now have the best of both worlds in a single device: OT for machine control and IT for data analytics.

Advanced HMIs: essential for a smart connected world

The IT and OT worlds have converged. IIoT-enabled business systems are closely integrated with the operations at the plant floor. Data is no longer just northbound. Business systems are supplied with near real-time data from production floors across the enterprise. Critical data on supply chain materials and movements can be pushed down to the plants in order to harmonize production and supply. This allows staff at all levels to have the same view into data that is necessary to optimize performance.

HMIs can also play dual roles. They are, of course, used in the traditional way to operate and monitor machines. They are also the instruments on the plant floor that relay collected data to the connected worker about current and predictive manufacturing events. The data can be visualized in dashboards via trend graphs, gauges, or geomaps. Connected workers can interface with the data by setting up their own metric and using ad hoc queries to get to the relevant information they need.

Most controls engineers agree that PLCs and edge controllers are optimized for deterministic sequential control. The IEC 61131 standard offers users multiple languages for machine control, such as Ladder Diagram, structured text, and function charts. These tools have been the proven workhorses of the control world for decades. They continue to be revered today, as they have evolved into effective methods for machine regulation.

As IT moves closer to the machine level, engineers need new tools to tackle plant floor data management, visualization, and network security application challenges. To meet these needs, the control industry is leveraging the benefits of Linux as an operating system for automation controllers. The openness and flexibility of Linux enables engineers to develop the applications they need. In many cases, designers can leverage open applications to help them meet their needs with a packaged application. Grafana is an example of an open Linux visualization and analytics platform that can be used within an HMI to effectively display complex metrics.

Today, HMIs must have the ability to be multitasking, providing industrial control as well as analytic and visualization tools. To meet these challenges, HMIs are equipped with traditional sequential control-based runtimes that operate on a Linux operating system. Controls engineers now have the best of both worlds in a single device: OT for machine control and IT for data analytics.

HMIs with IIoT connection tools

Empowering connected people with information is critical in this globally competitive industry. In order to transport data between IIoT cloud-based systems, plant floor HMI and other control devices are necessities for smart manufacturing.

IIoT cloud systems have multiple formats for the dynamic exchange of data. However, despite these alternatives, Ethernet-based MQTT and OPC UA protocols continue to be widely used. These protocols give engineers the ability to transmit complex data in a very efficient and familiar format. HMIs for the IIoT-connected world will need to be designed to easily and securely exchange information with cloud-based applications using standard protocols such as these.

Integrated HMI and PLC

Advances in technology continue to add features and reduce costs for the devices we use today. This is also true for HMIs. Industrial control manufacturers leverage evolving technology to enhance features, decrease electronic componentry space requirements, and increase processor power all while reducing costs.

Today, integrating your HMI and PLC into one device makes sense. The basic HMI of the past did a great job of replacing physical switches wired to a PLC, saving hardware and wiring costs. Advanced HMIs now are going a step further by taking on the functions of the PLC, helping to reduce duplication. Cost savings can be gained by having one device do the job of multiple devices. Engineers can use one software tool to develop both control logic and HMI graphic screens. There is no need to spend time importing and exporting tag databases between different applications or storing and maintaining two sets of application files. Integrating the HMI and PLC into one high-performance device will help reduce the cost of ownership of any machine.

Modern HMIs

The HMIs of today are not your father's HMI. Modern HMI needs to have the power, openness, and flexibility demanded by the complex applications of this smart connected world. End users continue to look for ways to empower their staff with real-time information while reducing the system cost of ownership. At the same time, engineers are demanding HMIs that can manage their traditional industrial control as well as run data analytics in parallel. It's time for everyone to start thinking of HMI as far more than a simple screen—this is a device that resides at the very backbone of the smart connected world.

Images courtesy of WAGO

ABOUT THE AUTHOR

Charlie Norz is the senior product manager for automation at WAGO Corporation, located in Germantown, Wisconsin. WAGO's parent company, located in Minden, Germany, pioneered CAGE CLAMP spring pressure connection technology and uses it in their extensive range of Interconnect, Interface, and Automation solutions.

ISA International Society of Automation
Setting the Standard for Automation™

isa.org/alc

AUTOMATION & LEADERSHIP Conference

4–6 October 2023 | Colorado Springs, CO, USA

The ISA Automation & Leadership Conference (ALC) is the automation event of the year—combining ISA's leadership conference with the best technical presentations from its automation conference series into an unparalleled event experience.

You will have the opportunity to network with ISA leaders and automation professionals from around the world—including managers, engineers, technicians, and subject matter experts. This conference includes:

- Society meetings and leadership presentations
- Vital technical sessions and access to an array of technical subject matter experts from the US, Canada, Middle East, Brazil, Malaysia, Spain, and India
- Interaction with other conference attendees—in person or virtually
- Exhibitor booths with insight and information on the latest products and technology in the industry
- The ISA Honors and Awards Gala

KEYNOTE SPEAKER

Mark Weatherford
CSO and SVP of Regulated Industries, AlertEnterprise and Chief Strategy Officer, National Cybersecurity Center

TRAINING OPPORTUNITIES

Training courses are under development and will be announced soon.

TO BE ANNOUNCED!

WHO SHOULD ATTEND?

- Automation Engineers
- Process Control Engineers
- Security Engineers
- QA Engineers
- Plant Engineers
- Manufacturing Engineers
- ICS Cybersecurity Engineers
- Digital Transformation Managers
- Engineering Managers
- Security Operations Center (SOC) Managers
- Compliance and Risk Managers
- Chief Information Officers (CIOs)
- Chief Information Security Officers (CISOs)

INTERESTED IN SPONSORING?

Check out our sponsorship prospectus at **isa.org/sponsor** or contact:

Kim Belinsky
+1 919-990-9404
kbelinsky@isa.org

Morgan Foor
+1 919-990-9267
mfoor@isa.org

Learn More

Using OEE Data to Improve Manufacturing Operations

By Kevin Welsch, SICK

As technology advances, plants reconsider their environments and adjust their processes to make the most of it.

As manufacturers continue to tackle the challenges of the modern operating environment, more of them are turning to data and technology to solve current problems and plan for the future.

A shortage of labor, the ongoing supply constraints, and high demand for products are pushing manufacturers to rethink how they run, monitor, and optimize their production environments. A place

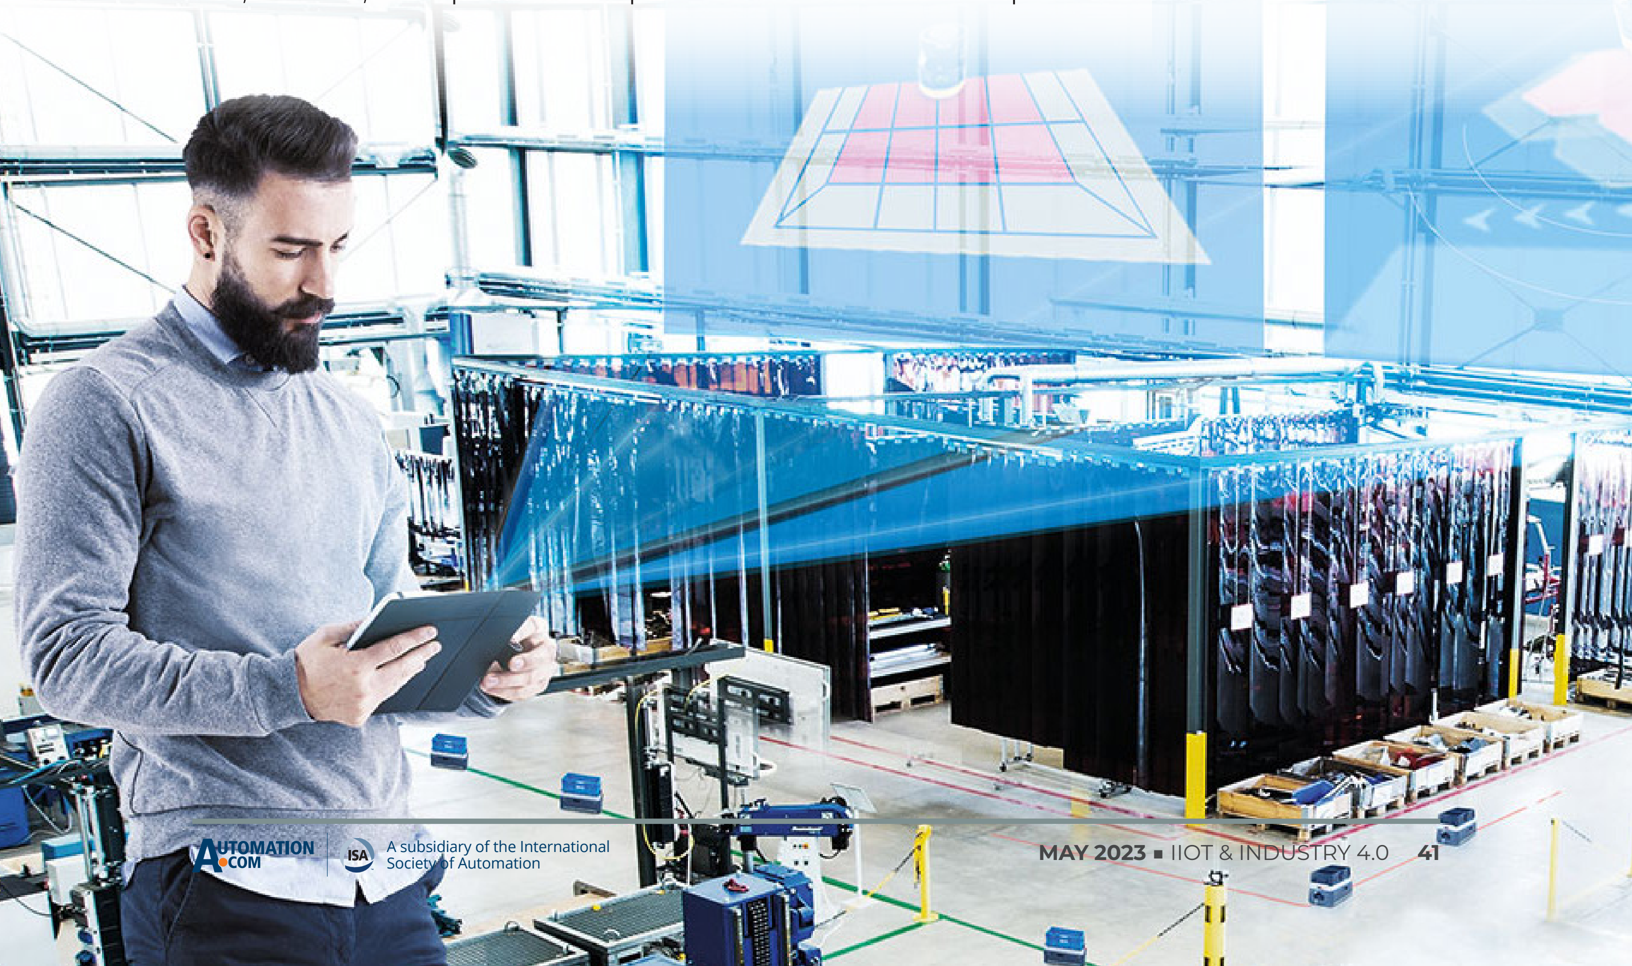

where seat-of-the-pants decisions and guesswork don't cut it, modern manufacturing requires a data-centric approach that helps companies identify problems, leverage opportunities, and pivot accordingly.

Checking all these boxes in the fast-paced manufacturing environment isn't easy, but technology has advanced to the point where it can help companies improve operations without having to add more labor, equipment, and production lines. Under constant pressure to do more with less—and with less downtime—companies are using data analytics to get more productivity out of their existing facilities in a safe and predictable manner.

OEE is a means to measure manufacturing productivity. It gives the information a manufacturer needs to be able to readily identify bottlenecks, points of failure, and areas of opportunity.

Of course, continuous manufacturing process improvement requires access to all data, not all of which is easy to capture, aggregate, or make actionable. Without this data, process inefficiencies, unexpected machine downtimes, and production losses emerge. By monitoring overall equipment effectiveness (OEE) and production lines, companies get all the information they need to be able to readily identify bottlenecks and points of failure, and to ensure process optimization across all production lines.

What Is OEE?

OEE is a means to measure manufacturing productivity. It helps to identify the manufacturing time that is truly productive and gives the information a manufacturer needs to be able to readily identify bottlenecks, points of failure, and areas of opportunity. With these data points at their fingertips, companies can optimize processes to meet the demands of today's operational environment.

A common metric used by manufacturers, OEE scales from zero to a 100% but is more than just one single number. Three different factors go into it: performance, availability, and quality. Together, these factors make up a manufacturer's OEE value. Leaders in their fields tend to run in the 85% range, while companies that are playing catchup usually have OEEs that range from 30% to 60%.

At its simplest, OEE measures manufacturing productivity and tells companies just how effectively they are (or aren't) utilizing their equipment. For organizations that invest much of their capital in equipment—and then don't use those assets to produce revenues or experience a lot of downtime on that equipment—low OEEs are fairly common.

A shortage of labor, ongoing supply constraints, and high demand for products push manufacturers to rethink how they run, monitor, and optimize their production environments.

A company might have a production line installed, but if this CapEx expense is only being used for three hours during a shift and only at half-speed, then the company is not using its assets effectively.

Keep the machines running

When equipment goes down on a production line, everything from inventory levels to supply chain management to customer service can be negatively impacted. Employees are left idle while the machines are fixed, throughput can come to a halt, and performance targets are missed. By Deloitte's estimates, manufacturing downtime costs companies about $50 billion annually, while poor maintenance strategies may hinder production capability by anywhere from 5% to 20% for a single plant.

Learning the cause of the downtime is the first step in reducing these numbers and keeping the machines running, and it starts with knowing the OEE of those assets. In most cases, the root of the problem is an electrical issue, a mechanical malfunction, or operator error. If it's an electrical issue created by an excessive number of loose wires in the machinery, for instance, then better training of electrical department employees might be in order.

With accurate insights into why the machines are failing (AKA fault metrics), companies can readily address the issues and minimize overall downtime. The problem is that not all companies have visibility into those fault metrics. Even for those that do, their OEE insight might be focused solely on the number, without really understanding why it continues to fall or what to do about it. Not much can be done with the OEE number without awareness of the areas that are in need of improvement.

Equipped with the right data collection and assessment tools, however, the same company gains high levels of visibility over performance problems, bottlenecks, throughput constraints, and other challenges. Using the data, manufacturers can effectively tackle these issues and increase their productivity and uptime. This, in turn, will result in a higher OEE number.

Using data to optimize productivity

A leading producer of sensors and sensor solutions for industrial automation applications, SICK makes sensors that are used on the shop floor to collect and analyze the data that manufacturers need to be able to run their operations at optimal levels. Using key data points like running status, product count, and quality counts, SICK's solutions encapsulate the insights in a software analytics platform that provides OEE calculations.

Manufacturers use the platform to see how well their lines are performing. With a few sensors and communication enabled by OEE analytics, manufacturers gain access to real-time and historical insights into meaningful KPIs that they can use to stabilize and/or optimize productivity.

SICK offers both a standard Package Analytics Platform and a Rapid Deployment Kit (RDK). Using either of these options, manufacturers can start with just one machine and then scale up to an entire production, packaging, or other line.

Companies can start small and zoom into the area of initial concern, knowing that the solution can then be expanded to other areas of the facility. By beginning at one end of the plant and making their way across the facility, companies can improve operations, performance, and quality as they go.

The OEE Rapid Deployment Kits combine software and hardware to help companies quickly begin using production data to drive better business decisions. The solution's OEE analytics and product lifecycle management (PLM) tools have been preconfigured and enable fast, easy commissioning and installation. They help companies drive improvement by providing a better understanding of production losses and reduce seemingly complex production problems

When one high-volume producer of pre-packaged gourmet foods noticed discrepancies in the volume between shifts, it realized that it needed a quick way to identify production metrics and bottlenecks. Implementing OEE concepts, it upgraded its production line and implemented a data analytics solution.

to simple, accessible information that manufacturers can use to improve efficiency and lower operating expenses.

Key metrics that the solution monitors include availability, shift duration, elapsed shift time, machine/system uptime, performance, expected and predicted outputs, actual outputs, and quality output. Using real-time monitoring of these and other critical OEE metrics, companies gain a better understanding of overall system behavior and trends, identify target areas for improvement, and make better, data-driven business decisions.

Equipped with these valuable insights, companies can reallocate labor to other value-added tasks, give their customer service and sales teams more accurate information, define roles for who can access the data, increase their sales revenues, and gain an edge on their competitors. They can also take a more proactive approach to machine maintenance and stability, which in turn supports higher equipment uptime.

Manufacturers are under pressure to produce more while also keeping systems running until they fail, but this is a flawed approach. One line that goes down can throw off a whole production schedule, which then pushes everyone into disaster recovery or reactive maintenance mode.

Don’t just go with your gut feeling

With an analytics-based OEE solution in place, manufacturers know how well they’re performing on all three metrics—availability, performance, and quality—and have an overarching view of how well they’re doing.

Based on that data, companies get a grade rating as they make improvements like mechanical adjustments or operational changes to standard operating procedures (SOPs). Then, they can use their OEE numbers to assess what is or isn’t working and make further adjustments as needed.

This data-based approach is much more effective than the “gut feeling” strategies that many manufacturers rely on to keep their lines up and running. Rather than just guessing at what might work, they can use the data to test out various continuous improvement initiatives on their lines and immediately see how those shifts impact performance.

Large food manufacturer runs smarter with OEE

When one high-volume producer and distributor of various local and authentic pre-packaged gourmet foods across 37 countries noticed discrepancies in the volume between shifts, it realized that it needed a quick way to identify production metrics and bottlenecks. Historically, it conducted manual process counts of the production lines between shifts—an approach that created numerous operational inefficiencies.

For example, manual entry errors in production quantities led to miscounts and inaccurate paperwork. These, in turn, decreased actual production number accuracy. The manufacturer also couldn't properly monitor inventory and production or predict production numbers with a high degree of accuracy. This resulted in over- and under- producing efforts, both of which impacted its profitability.

Combined, these inefficiencies were affecting the company's daily production and impeding its ability to meet growing customer demand. Working with SICK, the manufacturer began implementing OEE concepts. It upgraded its production line and implemented a data analytics solution, the latter of which would enable better access to data and insights into how to best reduce risks and challenges.

Using SICK's Rapid Deployment Kit, the food manufacturer gained a better understanding of production losses. The encoder ensures that the machine is running and sends a signal to the SIG200 to confirm this. Then, the product is counted using presence sensors as it runs down multiple production lines simultaneously. Finally, the machine captures product data plus any missing pieces or flaws in the production line and sends that information up to the SIG200 for review.

Once the goods pass quality inspection, the finished products are stacked for processing and shipping. With this solution in place, the manufacturer now has a baseline and can track real-time and historical

OEE metrics over time, leading to better overall machine availability and a higher quality product to offer its customers.

Manufacturers do more with less

A key metric for measuring manufacturing productivity, OEE helps companies identify losses, measure progress, eliminate waste, and improve manufacturing equipment productivity. And while the OEE metric itself has been around and in use for some time now, for the most part companies have been tracking it on paper and Excel spreadsheets. Both approaches are prone to data entry errors in an environment where three-second "micro stops" that repeat themselves may significantly impede overall performance.

Using sensors, data, analytics, and a user-friendly dashboard, SICK makes it easy for manufacturers to measure OEE, address issues, leverage opportunities, and implement a culture of continuous improvement. These are key wins in a business environment where all companies are under pressure to maximize throughput and do more with less, and shutting down a line to install new software or hardware on it isn't a workable option.

The RDK can be deployed even while the line is running, which is important for manufacturers that don't have the luxury of shutting down their operations. With RDK, the system is up and running, and gathering the data quickly. This immediately shows exactly why those micro stops or other issues are occurring so manufacturers can take the steps necessary to improve OEE across all three metrics—availability, performance, and quality.

Images courtesy of SICK

ABOUT THE AUTHOR

Kevin Welsch is an industry marketing manager and acts as the marketing liaison for SICK's Factory Automation Regional Sales Team, Indirect Distribution Channel with a focus on consumer goods, machine builders, electronics, and solar business initiatives. Kevin is a seasoned advertising, marketing, and sales professional within the industrial automation industry.


For MQTT, Smarter is better!
DataHub®
DataHub® is the only broker that parses and manages MQTT data intelligently. Data coming from multiple devices can be filtered, aggregated and processed, monitored and secured while ensuring data consistency from the field to the dashboard.
If you're using MQTT for IoT, be smart! Get the DataHub Smart MQTT Broker.
Learn more about DataHub®
SKKYNET™
SECURE INDUSTRIAL IoT REDEFINED
DataHub® is a registered trademark of Real Innovations International LLC, used under license

Growing IIoT Challenges **Require Smarter MQTT**

Larger, more complex systems can benefit from a smart MQTT broker.

By Xavier Mesrobian, Skkynet

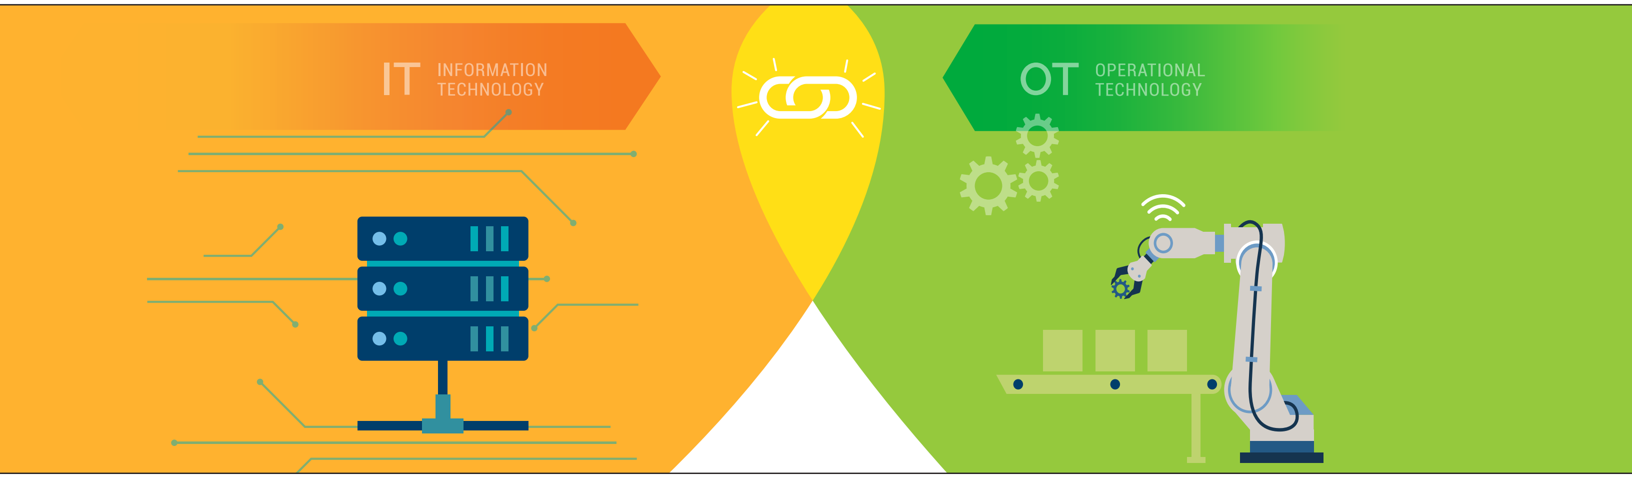

Since the introduction of Industrial IoT (IIoT) and Industry 4.0, we've seen an upsurge in interest in the MQTT protocol. Initially developed as a way to send data from field devices to a central location, MQTT is efficient, quick, and secure. It seems ideally positioned for connecting operational technology (OT) systems of all types to corporate information technology (IT) systems and the cloud.

Yet, as IIoT applications grow in size and complexity and projects move from pilot studies to large-scale, enterprise-wide applications, engineers are counting on MQTT to connect not only sensors and actuators in the field. Edge devices, SCADA systems, IoT gateways, and more are being linked to various tools on the IT side—historians, data lakes, AI engines, and other analytical instruments. This broader range of applications presents challenges to the MQTT protocol that was intentionally kept simple to ensure speed and flexibility.

Now, instead of each connection carrying data from a single device, MQTT is being called on to send collections of data values. Where once all devices may have been identical, today's complex systems often connect a variety of devices over different data formats. The simple, direct security model of device-to-client is not sufficient anymore when networks need to be isolated using DMZs, requiring multiple-hop connections. A new specification, Sparkplug B, was introduced to address some of these challenges, and yet there are ways that it, too, can be enhanced.

Get smarter

These challenges demand that MQTT get smarter. By design, MQTT is a transport protocol, like a postal service carrying letters. The service doesn't know or care what's in the letters. But what if we make the MQTT broker smart? What if we give it the ability to read and understand the messages it carries? It could then parse them and handle them more intelligently. And what if the broker could communicate with the senders and receivers themselves? It could then inform them of network status or which clients might have disconnected.

This kind of smart broker would be invaluable for the growing demands being put on MQTT. Let's look in more detail at what's needed and how making MQTT smarter can help it meet the challenges of IIoT and Industry 4.0.

Data collection

For many IoT and OT-to-IT applications, the simple device-to-broker MQTT connection is not sufficient. On large-scale systems with hundreds or thousands of connected devices, the data streams might need to be consolidated into a few or even one MQTT connection. This is particularly true for cloud services that accept only one client connection or that charge on a per-connection basis. And in many scenarios, MQTT is being used alongside other industrial protocols, such as OPC UA.

A smart MQTT broker can collect and aggregate incoming data messages, parse them, and translate various MQTT message types into one. And if it can read data in other common protocols like OPC or Modbus, it is not too much of a stretch to be able to convert that data into the same MQTT message type as well.

Data consistency

In a real-time industrial system, data consistency is critical. An operator monitoring an HMI or SCADA system needs to know exactly what's happening on the physical device. Data that's stale or out of correct time sequence can lead to incorrect decisions. Also, any disconnects or network irregularities must be known. A smart MQTT broker leverages its ability to parse messages, along with smart message queueing, to ensure data consistency.

- **Smart message queueing** – Real-time systems need smart message queueing in order to handle message overload. This happens when a data producer, like a sensor or other device, sends data faster than a consumer can receive it. A chronic overload requires the broker to drop messages. A smart MQTT broker can implement an intelligent

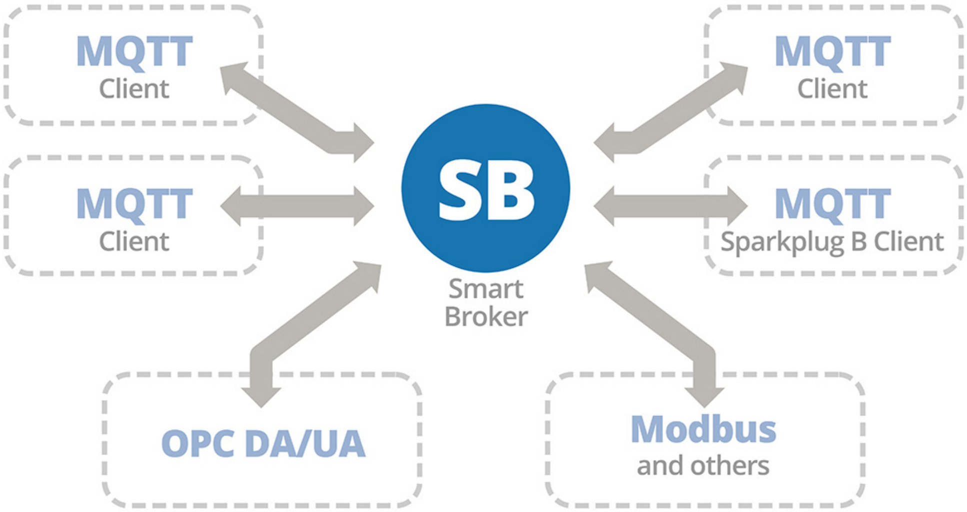

A smart MQTT broker can collect and aggregate incoming data messages from MQTT and other protocols, parse them, and translate them into one MQTT stream.

message queue that examines the message content and ensures that the latest value of every data item is delivered, even when earlier values are dropped. This keeps data at the consumer consistent with the physical reality of the producer.

- **Latest value** – Having the very latest value of the data is critical in an industrial system. Suppose, for example, in a burst of activity a pump is switched on and off many times, with the final position being “OFF.” If that final MQTT message gets dropped by the broker, the HMI or SCADA system will show the pump as “ON.” This kind of inconsistent data can lead to costly errors and system malfunctions. A regular MQTT broker without smart message queueing may drop that final, latest value, whereas a broker with smart message queueing ensures that it gets delivered.

- **Time order** – Time order is preserved in a single MQTT message topic, but not necessarily among multiple topics. Events coming from different devices that occur in the order A then B then C could be delivered to an application as C then B then A, or any other ordering, which is an error in many industrial-control use cases. A smart broker can preserve time order as it converts messages to other protocols for transmission to control systems or retransmission across a network.

- **Connection status** – Regular MQTT brokers do not have a way to indicate that a data source is disconnected. The consuming application cannot tell the difference between an old value from a sensor that has failed or a current value that has simply not changed recently. The “last will” mechanism in MQTT designed to deal with this requires unreasonable levels of coupling between the producers and consumers of data, resulting in duplicate configuration and increased integration and maintenance costs.

A smart broker that monitors the condition of the data producers and the network can assign a quality code to each message and update it with each value change. This information can be included in the outgoing MQTT message. As a result, data consumers have some way to tell why a value is not changing.

Data security

Industry security experts and government agencies recommend isolating networks for connecting OT and IT systems. The preferred approach is by using a DMZ. NIST document SP-800-82 sums up it up like this: “The most secure, manageable, and scalable control network and corporate network segregation architectures are typically based on a system with at least three zones, incorporating one or more DMZs.”

These three zones are the control zone (OT), the corporate zone (IT), and the DMZ in the middle. Using a DMZ ensures that there is no direct link between corporate networks and control networks, and that only known and authenticated actors can enter the system at all. The SP-800-82 document describes the value and use of firewalls to separate these zones and to ensure that only the correct data passes from one to the other.

Multi-hop daisy chain

Implementing data flow through a DMZ is problematic for MQTT, as this kind of connection typically requires two or more servers linked together one after the other in a daisy chain. The Quality of Service (QoS) guarantees in MQTT cannot propagate through the chain, making data at the ends of the chain unreliable.

One reliable solution is to convert the MQTT message into a different format that can be passed over the network from server to

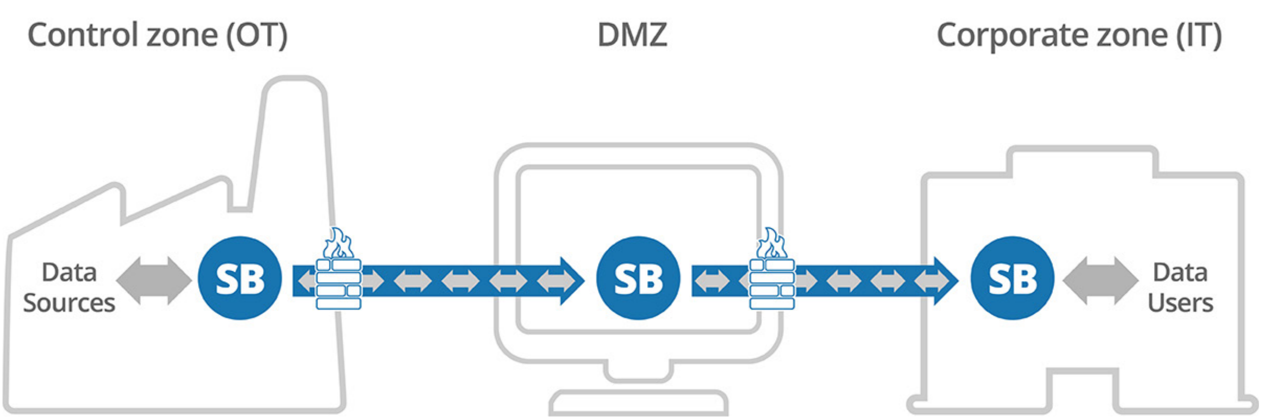

Industry security experts and government agencies recommend isolating networks for connecting OT and IT systems. The preferred approach is by using a DMZ.

server until it reaches its destination. The device producing the MQTT data would be connected to an instance of a smart broker. The broker, capable of doing data conversions, would pass the data, along with its quality information, via a secure protocol to a second instance of the smart broker, which would convert the data back into MQTT.

Ideally, the protocol used would offer SSL encryption, preferably with support for the most recent versions, such as TLS 1.2 and TLS 1.3, as well as use and enforce server certificates. Also, the smart broker should be able to replicate the ability of an MQTT client to send data outbound from a firewall without opening any inbound ports. It is critical that this valuable security feature of MQTT be retained.

A smart MQTT broker can collect and aggregate incoming data messages, parse them, and translate various MQTT message types into one.

Sparkplug B enhancements

The Sparkplug B specification for MQTT was introduced to resolve interoperability issues between vendors by defining how data is sent and received. Sparkplug B classifies MQTT clients as either edge of network (EoN) devices that produce data or as applications that consume data. Each Sparkplug B device produces messages of various kinds, such as a BIRTH message to show that it has come online, DATA messages for sending data, or a DEATH message when it goes offline. Any Sparkplug B application that is online receives these messages and is thus kept informed of which data is coming from which device.

All of the smart broker capabilities discussed so far apply to a Sparkplug B-based system. Additionally, a smart MQTT broker may provide other features to further enhance Sparkplug B connectivity.

- **Synchronizing all applications** – Because it is aware of all connections, a smart broker can synthesize a BIRTH message for each connected device whenever a new application comes online. This allows that application to receive DATA messages from all currently connected devices, eliminating issues related to start-up order. The recent Sparkplug 3.0 specification adds a mechanism for an application to announce its presence such that EoN devices will transmit BIRTH messages, eliminating the need for this function in the broker.

- **Responding to errors** – In addition to its ability to identify out-of-order or lost MQTT messages, a smart broker should also be able to automatically disconnect a Sparkplug B device when these kinds of errors occur, causing it to reconnect. This would cause the device to re-send its BIRTH (start-up) message, which will resynchronize all receiving applications, thus maintaining a single version of the truth.
- **Resolving failed writes to devices** – Another useful feature would be to monitor all write requests from applications to devices to ensure that the specified data value was written on the device. If the smart broker detects that the value on the device did not change, it would force the device to disconnect, causing it to retransmit its BIRTH message. This would resynchronize all applications listening to that device and is another way to maintain a single version of the truth.
- **Adding data quality information** – For systems that need to convert Sparkplug B data to other protocols, a smart broker could add quality information. For example, when converting Sparkplug B data to OPC, it could add OPC data quality. When a Sparkplug B connection is lost, the smart broker can update the data qualities of all related OPC items, alerting downstream applications to the loss of connectivity.

A smarter future

As valuable as MQTT is for device-to-server data communication, it needs to get smarter to take on current and future challenges of OT/IT, Industry 4.0, and Industrial IoT. A smart MQTT broker can collect data from multiple incoming message types, and even other protocols. It can ensure data consistency over the entire path of the message from data producer to data consumer, where the consumer always has the latest value with an indicator of data quality. A well-designed smart broker can also be used to securely connect MQTT data producers and consumers across DMZs and other multi-hop network configurations. These advantages and more can strengthen Sparkplug B implementations as well. For today's challenges and those that lie ahead, MQTT must get smarter.

Images courtesy of Skkynet

ABOUT THE AUTHOR

Xavier Mesrobian is the vice president, sales and marketing, at Skkynet Cloud Systems, which allows companies to securely acquire, monitor, control, visualize, and consolidate live process data in-plant or over insecure networks with no VPNs or changes to IT policy required.

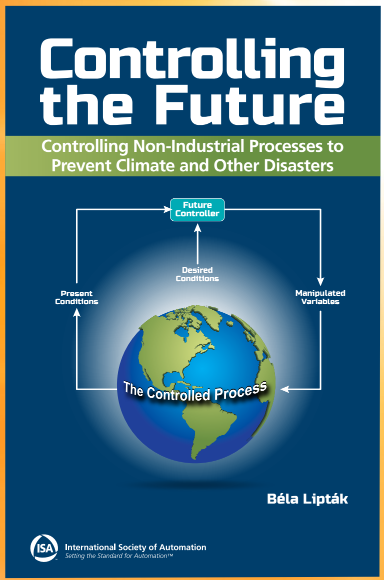

The Data Mesh—An Advanced Distributed Data Lake Architecture

By Jim Redman, ErgoTech Systems, Inc.

Rapid advances in machine learning (ML) and artificial intelligence (AI) are extending the range of problems that can be solved and making building and training models easier. The primary barrier to advances in this field in manufacturing continues to be incomplete or inaccurate data sets.

In the March issue of Automation.com, we discussed the traditional centralized data lake architecture and the benefits of moving to a distributed data lake. We're taking that one step further to create a flexible, high-performance, highly scalable data integration platform for manufacturing—a true "data mesh."

In this architecture, data owners, individuals or teams with in-depth knowledge and understanding of specific data sets, are in charge of creating, processing, and providing their specialized data. This distributed,

> Treating data as a product puts responsibility for its accuracy in the hands of data owners and spurs collaboration.

decentralized, loosely connected system treats data as a product, allowing the data owners, who have the most expertise and insight into the data's importance, to assume responsibility for its management. This approach empowers them to collect and publish data without relying on a centralized team, which leads to a more complete data set. Data consumers, such as computer and data scientists, are then free to discover and use that information for AI/ML or other applications.

Rather than try to copy all our existing sources of data into a central system, we need to "wrap-and-embrace" diverse data sources; i.e., integrate these distributed platforms so that they can be searched as a single platform.

Data collection in the AI era

It's clear that AI/ML is transformative. The substantial advancements in large language models, like ChatGPT, represent just one of numerous AI/ML techniques that have already begun to transform all industries, encompassing internet search, software development, and document generation. In the manufacturing sector, AI-driven advancements can optimize production processes, reduce waste, and improve quality control, opening up new possibilities for efficiency and innovation. Despite the impressive capabilities that AI/ML approaches currently demonstrate, they are still in their early stages of development and have immense potential for growth and improvement. In a detailed review of the semiconductor industry, McKinsey estimates that "AI/ML now contributes between $5 billion and $8 billion annually to earnings before interest and taxes," but also that this "reflects only about 10 percent of AI/ML's full potential within the industry."

Artificial intelligence and machine learning applications have an immense appetite for data, and suboptimal results can stem from incomplete or inaccurate data input. Models such as ChatGPT are trained using vast amounts of data sourced from the internet. In the context of manufacturing, it is crucial to gather data that is specifically tailored to our environment and process to achieve the desired outcomes.

Existing approaches—data warehouses and monolithic data lakes

Data warehouse

Manufacturing generates vast amounts of data from diverse sources such as machinery, sensors, and enterprise systems. Historically, data management in this sector has relied on "data warehouses," which are often built on sizable relational databases. The objective of a data warehouse is to establish a centralized storage space for data, generally used for reporting and data analysis purposes. Data warehouses are adept at handling traditional "structured" manufacturing data, such as values obtained from PLCs or time-series data, which can be easily organized and comprehended.

As AI/ML models continue to advance, they are now capable of deriving valuable insights not only from traditional structured data but also from unstructured data. Unlike structured data, unstructured data lacks a specific format or organization, presenting a more complex landscape for analysis. Examples include text documents, waveforms, log files, blueprints, schematics, images, and videos. It is characterized by a high degree of variability and a lack of predictability.

Unstructured data offers engineers and data scientist the opportunity to gain significant insights that cannot be obtained solely from conventional, structured manufacturing data, thus delivering additional context and understanding. Reports and Markets suggests that "over 40% of the operational value of IoT is extracting and monetizing unstructured data."

Data warehouses are poorly equipped to manage this unstructured data. Transformations to coerce it into a more structured format suitable for a data warehouse can limit the insights that organizations can derive from unstructured data, as they may not have the tools to perform advanced analytics, such as natural language processing and image recognition. Data lakes aim to overcome this limitation by using alternative data storage and processing solutions that are better suited for handling unstructured data, such as NoSQL databases and big data processing platforms like Hadoop and Spark, and by storing it in raw format or with very little modification.

In both the monolithic data lake and data warehouse, the goal is to consolidate everything in a single storage system. It's immediately obvious that it's wasteful, both in terms of storage and data collection resources, as well as data management effort. It's difficult to see a purpose for duplicating data that is already part of a working system, such as a manufacturing database, MES, ERP, or similar.

Few organizations of any size have a single data warehouse, database, or manufacturing system. Different departments and divisions build data storage for different use cases, or IT systems are adopted as part of an acquisition. These "silos" of information are not readily integrated, often having been created for different functionality and being managed by different organizational units. This is problematic for any monolithic approach, data warehouse or data lake.

Intuitively, the architecture does not match the nature of data sources in manufacturing. The number of data sources in manufacturing is huge and distributed—and all the sources produce data in different formats. Every piece of equipment, PLC, subsystem, and even smart breakers and small sensors are generating data that's potentially valuable for AI/ML applications. The collection and management architecture should match the real nature of the data—distributed and very variable.

All of these data sources have internal domain experts, someone who really knows the tool or application and how and what data to extract These domain experts should be empowered to manage their data "product" to easily add, remove, or upgrade sources.

Data mesh

What we need is not to put all this data into a centralized system, but to be able to query, view, and extract this data **as if** it were in a single system. Rather than try to copy all our existing sources of data into a central system, we need to "wrap-and-embrace" diverse data sources; i.e., integrate these distributed platforms so that they can be searched as a single platform.

With this concept, we move from a centralized system to a heterogeneous and distributed set of data platforms. This is possible with a distributed architecture using an industrial internet of things (IIoT) or edge approach whereby small, self-contained applications ("microservices") close to the equipment can be managed by the data owner. These platforms are integrated by "data virtualization," allowing users and applications to query the data without caring about where or how the data is stored—the "data mesh."

Data meshes focus on treating data as a product and emphasize the importance of cross-functional, domain-oriented teams that are responsible for their own data. The main principles of data meshes are:

Domain-oriented ownership: Data is treated as a product, with individuals or teams taking ownership of their data and its quality.

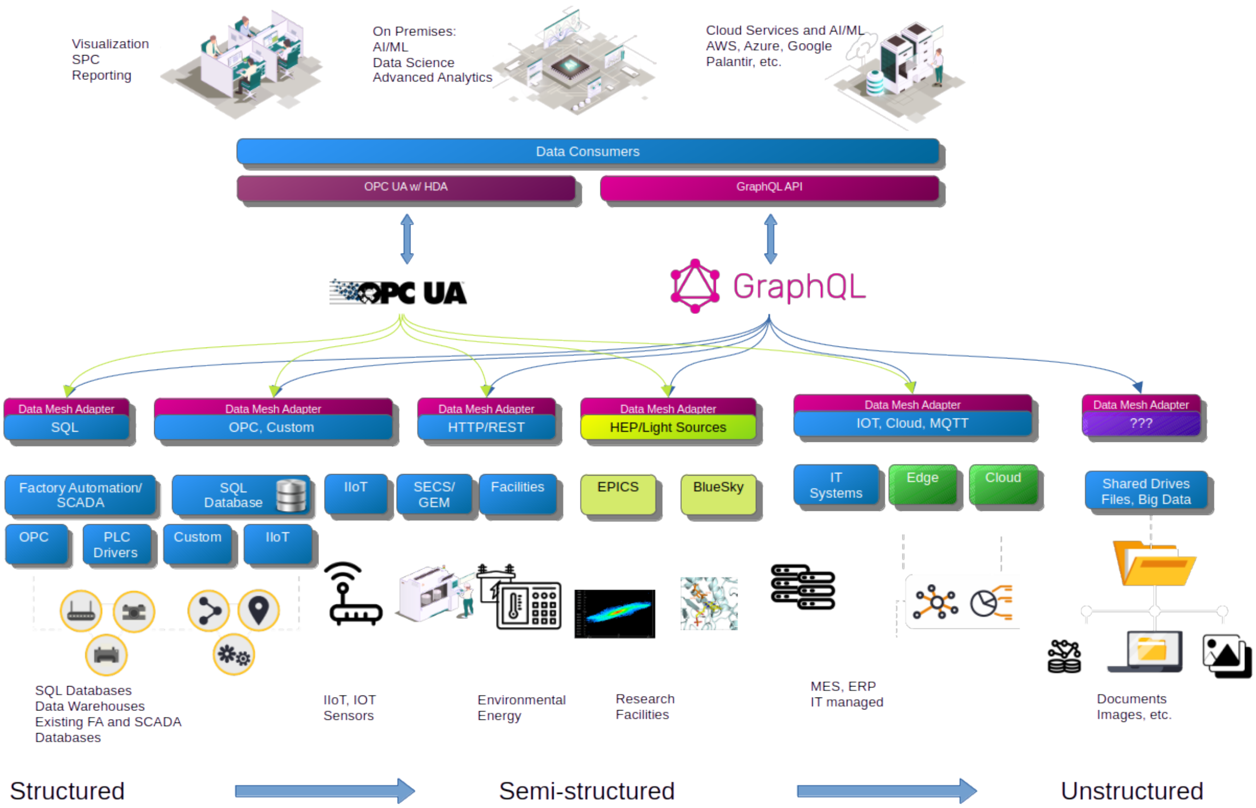

Data mesh: Harnessing varied data for enhanced AI/ML analytics, productivity, and quality.

These teams are responsible for generating, processing, and serving their domain-specific data.

Self-serve data infrastructure: Data meshes encourage building a self-serve infrastructure platform to support domain teams. This platform should be easily accessible and empower teams to discover, access, and use data without relying on a centralized team.

Product thinking for data: Data is treated as a product with its own lifecycle, from creation to usage. This mindset helps organizations focus on the value of data and its usability for AI/ML or other consumers.

All existing data platforms, the data warehouses, databases, MES, ERP, IT systems, and shared drives, become a part of the data mesh, not by moving the data from these systems to a central system, but by providing a virtualization layer.

The approach is also hugely scalable. As the volume, variety, and velocity of data grow, a data mesh can scale to handle large amounts of information across multiple sources without compromising performance or user experience.

Solutions for data ingestion and integration

The data mesh revolutionizes data management and offers numerous benefits, including increased scalability, flexibility, and decentralization; however, it must address the issue of data ingestion and seamlessly integrating disparate data sources. The data owners understand the data but are typically not trained software developers and so need easy-to-use tools to construct efficient ingestion pipelines. A data mesh must also provide a comprehensive solution for seamless data fusion and an intuitive and streamlined platform for data access.

Tools to empower data owners

Just as with a traditional data lake, the data mesh requires an ingestion pipeline. In a monolithic data lake, this is typically a software application created by the IT department. The complexity and management of this application increase as the amount and type of data being added to the data lake increase.

In a data mesh, this becomes much simpler. "Ingestion pipeline" now really just means "describe the data." This task has also now moved from IT to the data owner or domain expert, and onto a single node. The process is hugely simplified, as we're asking the domain expert to add just the data they know well without needing knowledge of any other data pipelines.

The remaining challenge is that the data mesh must provide user-friendly tools to allow non-programmers to easily define data, both structured PLC, sensor, timeseries, etc., and also unstructured images, spreadsheets, text documents, and so on. A new class of drag-and-drop tools is emerging that allows not only this ingestion but also the ability to export data in any format required by its consumers.

Client access should be abstracted from the data locations, and queries should be focused on what data is required, not how to access it; in other words, data virtualization.

These solutions empower data owners to streamline the process of creating ingestion pipelines, simplifying their tasks and allowing them to unlock the full potential of the data mesh infrastructure.

Data virtualization

For the data mesh to be a practical solution, distributing data must not unduly increase the complexity for the consumers of the data—the clients. Client access should be abstracted from the data locations, and queries should be focused on what data is required, not how to access it; in other words, data virtualization. Approaches to data virtualization, commonly used with SQL databases and data warehouses, include a REST API, MQTT, and many other IoT protocols. While these all have value, and should be supported by a data mesh, the primary limitation of these approaches is that access is largely defined by the creator of the interface, not the user of the data. The data set available is rigid, and the client may receive more data than required (over-fetching) or need to make multiple calls to access and combine the required data (under-fetching).

Additionally, the data virtualization layer should be mostly invisible during the data ingestion process. Data owners should not be required to possess expertise in any virtualization technology; the data mesh software should make the information available in the necessary format, making the process seamless for the data owner.

For conventional manufacturing data (structured data), OPC UA, with support from small systems to the cloud, can be a valuable virtualization layer, but it is not well suited for unstructured data.

In 2012, to address the limitations and inefficiencies of their existing REST APIs, Facebook (now Meta) developed GraphQL. GraphQL was open-sourced in 2015 and has been widely adopted and used by many organizations and companies, as well as being a standard for developing APIs. GraphQL is a query language and runtime for building and executing client-server queries.

Distributing and federating data has the potential to improve performance by allowing the execution of complex queries to span multiple data sources and enabling each service to resolve its part of the query.

With GraphQL, the client makes a request, a query, specifying the fields it wants to retrieve. The server responds with the requested data. This allows the client to retrieve exactly the data needed, resolving the problem of under- and over-fetching. GraphQL supports data of all types, whether it be structured or unstructured.

GraphQL allows multiple services to be combined into a single, unified system. This “federation” is an essential concept for a scalable data mesh, providing the ability to independently create data sources while still delivering a consistent and unified API for clients to consume.

Distributing and federating the data in this fashion has the potential to improve performance by allowing the execution of complex queries to span multiple data sources and enabling each service to resolve its part of the query, minimizing unnecessary data transfer, and leading to faster response times and reduced bandwidth usage.

Data mesh: empowering manufacturing

Implementing a data mesh in manufacturing holds the potential to significantly enhance operational efficiency and innovation. Utilizing a decentralized, domain-driven architecture, the data mesh facilitates seamless data discovery, encourages collaboration, and enables organizations to make real-time data-driven decisions to fully realize the potential of their digital manufacturing initiatives.

A primary advantage of the data mesh approach is that it empowers data owners to manage and provide the data themselves. This helps break down data silos, promote a culture of data-driven decision-making, and develop the notion of data as a product. As a result, organizations can foster innovation, increase agility, and sustain a competitive edge in an industry that is becoming more complex and fast-paced.

Images courtesy of ErgoTech

ABOUT THE AUTHOR

Jim Redman, as president of ErgoTech Systems, Inc., was delivering what has become "IIoT" systems way back in 1998. ErgoTech's MIStudio suite reflects his holistic vision to provide a single tool for integration and visualization from sensor to AI, and from tiny IIoT to worldwide cloud.