Living on the Edge:

HOW COMPANIES CAN BUILD AND OPTIMIZE INDUSTRIAL IOT INFRASTRUCTURE USING DATA FROM THE EDGE

The edge of a network, where the worlds of Operations Technology (OT) and Information Technology (IT) meet, is the starting point on an organization's journey to implementing Industrial IoT.

"IoT represents a democratization of existing technologies. The key challenge lies in organizing these technologies into new structures while removing time and distance barriers between the edge and the enterprise. Information and decisions will suddenly engage the organization more quickly and more broadly. When information and decisions are focused on the right business problems and engage the proper parties, they will deliver tremendous value. When they are not applied properly, they will cause tremendous disruption."

(Gartner Planning Guide, October, 2016)

Today's Industrial Revolution, known by many titles including the Industrial Internet of Things, the Industrial Internet, Made in China 2025, or Industry 4.0, is powered by data from the "things" that fill our factories and supply chains. These things include sensors, actuators, robotics, CNC machines, programmable logic controllers (PLCs), programmable automation controllers (PACs), and all the other assets that collectively enable the manufacturing of products and services. In essence, in the industrial IoT, "dumb" devices become "smart", or more accurately, "connected". They possess varying levels of processing functionality, ranging from simple sensing and actuating, to control, optimization, and full autonomous operation. Connecting devices to a network where they can exchange more substantial amounts of information and engage with each other is only the first step. The real value lies in the data that is generated from orchestrating relationships between the sensors, controllers, edge-computing nodes, devices and their environment, the analytics that can be performed with the data, and the insights into efficiency, yields, and maintenance that can be discerned through the analytics. Data, placed in context and properly synthesized, is at the heart of every industrial IoT system.

In the last decade, the rapid evolution of microprocessor technology in processing power and functional integration has resulted in technologies like system on a chip (SoC) and micro-electromechanical systems (MEMS) that enable micron scale integration of complex system functionality while greatly reducing power, footprint, and cost. With these new capabilities in hand, engineers have been able to create increasingly complex control systems and increasing levels of automation and integration on the factory floor. It is forecast between 26-30 billion devices will be connected to the internet by 2020, creating an economy of over $3 trillion in new business opportunities with the large majority of this increase from the industrial sector (*"The Internet of Things: Sizing Up the Opportunity," McKinsey*). All of this activity is expected to generate a flood of data of almost unimaginable proportion.

Advances in communications and the evolution to virtualized hardware coupled with the messaging infrastructure of internet technologies – all driven by a singular focus on open standards – have provided the infrastructure we now know as cloud computing. Cloud computing fundamentally changes the feasibility of building an industrial IoT business platform and the resources and cost required to get it done. Organizations can replace the up-front fixed expense with an infrastructure that accepts the data produced by automated manufacturing systems with low variable costs that scale with their business. Using internet technologies, organizations can deliver results faster by spinning up hundreds or thousands of servers in minutes, eliminating the need to plan for and procure servers and other IT infrastructure weeks or months in advance. Plus, these new cloud systems provide reliable, secure networks that can dynamically scale as circumstances dictate and can easily integrate new applications and services that analyze this data and provide new insights. Most importantly, enterprises only pay for the services they use, as they use them.

The availability of new, integrated sensor data, combined with advancements in connectivity, security, interoperability, and analytics, creates immense potential for organizations. It enables better monitoring, information gathering, role-based information presentation, situational awareness, process optimization, and predictive maintenance for operators. Organizations will quickly realize the benefits of better yields, more productive assets, improved safety, and reduced costs. The flood of market studies and analyst reports over the last couple of years certainly supports the opportunities IoT presents for transforming business. Based on current usage trends and the rate of IoT adoption, industry analysts predict that by 2020 more than 43 trillion gigabytes of data including audit logs, inventory, asset performance, work instructions, and yields will be generated by connected devices per month that will need to be processed in cloud data centers. (*Gartner Press Release, https://www.gartner.com/newsroom/id/3598917*)

So, What about the Edge?

In the industrial IoT, the edge of the network is that place where the Operations Technology (OT) domain, characterized by manufacturing plants, resource extraction sites, and the transport systems to move everything from raw materials to finished product – in short, all of an organization's assets used to create goods and services - bumps up against the traditional Information Technology (IT) domain. The edge typically consists of sensors, controllers, actuators, tag readers, communication components, gateways, and the physical machines themselves. The edge is where operational components connect, communicate, and interact with each other, with SCADA systems, historians, manufacturing execution systems (MES), and in some cases, directly with components in other locations or edges. The edge can be as small as a single, physical device with a direct connection to an IoT platform, or as large as an entire manufacturing facility with all of its equipment and processes, or anything in between.

In order to enable new business insights through the analysis IoT provides, data generated by OT monitoring and control systems must be integrated and accessible to the digital world of IT systems data processing. Once ingested by the tools of IT, this data can be normalized, analyzed through a wide variety of analytical methods, and converted into actionable insights. Such valuable insights can be leveraged to reduce asset downtime, improve inventory management, and increase productivity.

Historically, there has been little interaction between the OT and IT worlds. Information technology – also known as the digital world - focuses on information processing, the supporting technologies and services, the communications infrastructure, and databases.

IT is focused around maintaining the financial, selling, and support activities of the products that operations produces. Standards organizations within the IT space are predominantly focused on openness, speed, and agility across all platforms, networks and users.

Operational technology directly monitors and/or controls physical devices, processes and events with the objective of industrial control and manufacturing automation. Standards organizations within the OT space focus on conformity, safety and suitability. The automation systems of the OT have been proprietary systems, isolated from the globally connected world of IT. In turn, OT and IT have evolved and matured with very different attributes. Consequently, neither environment was designed to work with the other. As a result, most companies are finding out that moving from industrial IoT strategy to value can be complex.

	Information Technology (IT)	Operational Technology (OT)
Purpose	• Support people • Process transactions • Provide information	• Control machines and processes
Performance Requirements	• High throughput demanded • Control physical world • Data confidentiality and integrity paramount • Fault tolerance less important • Non-real time, consistent response	• Modest throughput acceptable • Manage Data • Human safety paramount • Fault tolerance essential • Real time, response is time-critical
System Operations	• Systems designed for use with typical operating systems • Upgrades and deployments straight-forward with automated tools	• Differing, proprietary operating systems, often with low security capabilities • Upgrades painful due to specialized control algorithms, modified hardware/software, and downtime
Architecture	• Enterprise wide infrastructure and applications	• Event-driven, real time embedded hardware and software
Connectivity & Communications	• IP-based corporate network • Standard communications protocols • Primarily wired with localized wireless	• Control networks, hard wired twisted pair • Many proprietary and standard protocols • Several including dedicated wire and wireless, radio and satellite
Ownership	• CIO and IT	• Engineers, technicians, operators, managers
Component Location	• Usually local, easy to access	• Can be isolated, remote and require extensive effort to access
Component Lifetime	• 3-5 Years	• 10-15 Years +

The Operations Side of the Edge and the Challenges of Implementing Industrial IoT

OT domain is comprised of an enormous array of tools, connectivity standards, and development and operating environments. It is a collection of technologies that includes standalone devices, embedded systems, and machine-to-machine communications across countless vertical applications. Industrial IoT designs encompass hundreds of proprietary and open source software platforms, purpose-built and application-specific, using a wide variety of methods to connect endpoints and applications.

Manufacturing automation is typically driven by a hierarchical structure based on the ISA-95 Standard. Data acquired by a sensor flows through an input/output, IO-module into a Programmable Logic Controller (PLC) or Micro-Controller (MC) which manages the local real-time control loop. Supervisory Control and Data Acquisition (SCADA) networks, Manufacturing Execution System (MES), and Enterprise Resource Planning (ERP) systems in the different layers of the automation pyramid govern the local control loops.

With a prime objective of reducing costs and improving margins that served as justification for their implementation, there is one attribute virtually all operations devices possess – the ability to continue to perform and to last for a long time, often in harsh conditions. Where average depreciation schedules on IT hardware and software tend to run five years or less, the expected lifespan of operations equipment spans at least 10 years.

Most of today's working devices like individual sensors and machines simply don't have the computing power required to process and filter the data they generate. At best they are pass-through devices with no intelligence. They interface through a bus or simple diagnostics interface and have no networking port. More intelligent devices like PLCs tend to focus on single-task automation functions, as they were selected to exactly match the requirements of the environment in which they operate. And with limited power and memory, fixed processing power, and size, OT devices generally cannot support new tasks through additional interfaces. They had not been specified or deployed with a vision to share the data they generate with other systems.

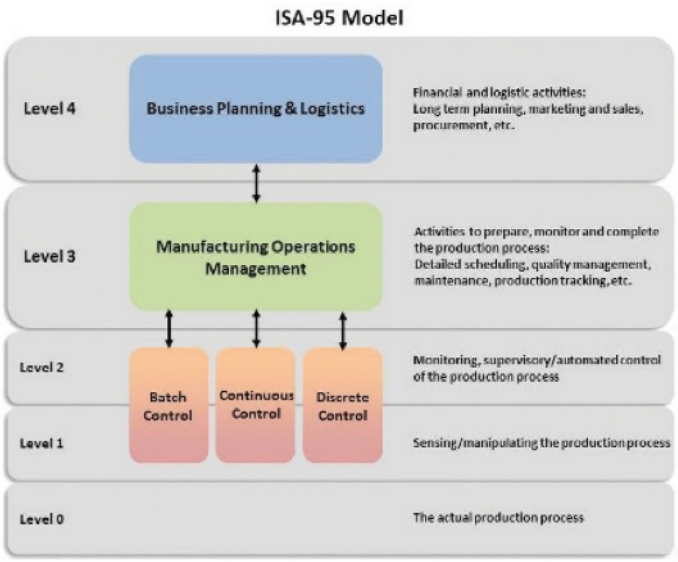

Many of the protocols used today in operations do not have the requisite security attributes for long-haul communication to the cloud. Companies prioritized real-time control for timely process data and safety-critical functions at the expense of long-haul information security. These protocols were deliberately stripped down and designed for low bandwidth networks when connectivity beyond the factory and the security to protect the network was not a concern.. Directly connecting devices that send unsecured data to the network presents a substantial threat plane that companies cannot ignore and must contain. Different strategies to isolate and protect these installations that do not involve making any changes to the devices themselves must be devised and implemented.

Quite often, existing OT telemetry networks have rather limited bandwidth for remote devices. Companies operating resource extraction sites might have thousands of legacy devices across multiple production sites with limited connectivity, power, and network bandwidth. As more devices are added, data production increases exponentially, requiring new levels of bandwidth, which can lead to service degradation, data latency, increased costs, and impacts on production.

There is one other attribute that virtually all sensors and devices possess - the capability to generate massive amounts of streaming data. The OT environment produces a very different kind of data than what most traditional IT organizations can handle – particularly time-series data like temperature

and pressure levels or KPIs - and at high rates of speed. Much of the data has a very short half-life. If it is not analyzed and acted on in near real-time, its value rapidly diminishes. Analytics and machine learning applications require high-frequency data in order to truly understand a process or machine. Failing to sample data at a high enough frequency could result in a completely different analysis or modeled behavior and potentially miss anomalies that could be significant. Moreover, high-frequency data generally translates into larger data sizes, pushing into the megabit and terabit range. This can easily exacerbate the bandwidth strain on an existing network infrastructure, leading to latency, degradation in service, or even outright loss of data. Sending all data to a remote processing center for time critical monitoring and analysis would not be an appropriate use of cloud computing if the data's usefulness is lost due to network latency.

Despite being capable of generating massive amounts of data, only a fraction is used by the PLC/MC/IPC to manage the control loop's automation task. Although the remaining data might not be relevant to the control process, it contains insights into a variety of other critical areas including crew behavior, lingering quality issues, and equipment health. As a result, the data collected becomes fragmented and siloed across the manufacturing domain, stored in various incompatible formats, and inaccessible to IT.

In turn, the situation many organizations face is reminiscent of the tale about the Blind Men and the Elephant. As the tale is told, a group of blind men who have never come across an elephant before try to learn and conceptualize what the elephant is by touching it. Each one feels a different part, but only one part, such as the side or the tusk. They then compare notes and learn that they are in complete disagreement. Eventually, they stop talking, start listening and collaborate to "see" the full elephant. The parable illustrates that one's subjective experience can be true, but that such experience is inherently limited by its failure to account for other truths or a totality of truth.

Like the blind men, with data siloed in different places throughout the organization, most data generated by factories is not currently used to generate actionable insight. According to an IDG Research Survey, over 90% of all data collected can go unused (IDG, December 2016). At best, data that does get analyzed is based on spreadsheet analytics. If someone wants to know the status of production, they either walk down the hallway covered with manually generated paper reports and graphs trying to discern if they are on plan, or they walk out onto the factory floor and take a look for themselves. Machine maintenance is usually based on historical data and requires careful downtime planning against production schedules. Invariably, the machines do not cooperate with these schedules, creating unplanned downtime with labor-intensive firefighting to bring the machine back on line. The same IDG Research Survey reported that 80% of senior executives looked forward to improved operational efficiencies and uptime as a top beneit of implementing industrial IoT - but 64% also felt that integrating data from disparate sources and formats in order to extract this business value is the single biggest challenge presented by the industrial IoT. Connecting to the devices that generate this data with all the challenges discussed so far and aggregating the data in a useful and cost effective manner is the first step in IoT implementation - whether from existing PLC, SCADA and Distributed Control System (DCS) systems, or from newly installed sensors.

The Edge, the Cloud, the Fog, the Dew, and Other Considerations

Today the terms edge computing and fog computing have evolved to become interchangeable, characterized by an organization's ability to choose where and how it wants to implement an industrial IoT solution without regard for the constraints of cloud-only architectures. And newer paradigms like dew computing will not only provide more alternatives to implement industrial IoT systems that fit the particular requirements of your operations, but will also deliver better toolsets to build, test, and implement them.

There is an unprecedented amount of activity tackling these challenges from a broad diversity of perspectives and establishing standards that can bring both the OT and IT worlds together cleanly and transparently. Some of the more prominent groups addressing the challenges of an all-inclusive IoT reference architecture include the Industrial Internet Consortium (IIC), the OPC Foundation, Smart Factory, and Industrie 4.0. Still, more traditional standards bodies like ISO, ANSI, EIC, Oasis Message Queue Telemetry Transport (MQTT), and the Open Connectivity Foundation Constrained Application Protocol (CoAP) focus on developing IoT standards in

a single part of the architecture, usually at the level of the OT devices or for a given industry.

One of the more interesting developments in the standards forum discussion is an evolution of industrial IoT architecture driven by groups including the OpenFog Consortium and the EdgeX Foundry. Rather than pushing all data to the cloud to provide the big data analytics that will drive new insights, there is a growing movement to push data collection and analytics closer to the OT edge.

The OpenFog architecture represents this shift from a reliance exclusively on cloud IoT models, to a new computational model that moves collection, aggregation, and analysis computation at or near the edge based on workload requirements and device capability. The segmentation of what tasks are done at the edge and what goes to the backend cloud are application-specific and could change dynamically, but it's no longer required that centralized decision-making occurs only in the cloud. Fog computing is designed to provide the ability to analyze data near the edge for improved efficiency (where delays are critical or there is limited bandwidth), or to operate while disconnected from a larger network (autonomy). Fog computing enables different fog instances to communicate with each other, enabling dynamic routing for resiliency and efficiency. (*Fog Consortium Architecture White Paper*)

The latest proposed paradigm for edge computing is the new concept of dew computing. Dew computing shares many similarities with fog computing – the ability to push data storage and analytics as close to the edge as possible, where it realizes the greatest benefit. However, where fog computing depends on a live, active connection to backend servers through a web browser, dew computing pushes rich microservices to edge computers that can run independently without a live permanent connection to cloud backend servers. However, when a dew computer is connected, it fully collaborates with the cloud service through synchronization, correlation, and other kinds of interoperation. Real-time data ingestion, analysis, and response can be provided in the most remote environments, particularly where communication services are intermittent or have failed. (*Yingwei Wang, Dew Computing Research, November 10, 2015)*

As both the cost and form factors of processors and memory continue to decline, the economics become much more attractive to move computing and data storage away from the cloud where enterprise-level applications and data reside. This certainly eases, if not solves a number of critical challenges businesses face in implementing industrial IoT. It enables businesses to distribute computing to the edge of the network through low-cost networks and hardened industrial PCs that can host localized and task-specific actions in near-real time while transmitting much less data back to the cloud. Additionally, this data can be transmitted with more modern protocols, like MQTT, specifically designed for efficiency and security.

Locating the server on the edge and connecting it directly to data sources helps to alleviate network bandwidth and reliability limitations while greatly reducing the security threat plane presented by current unsecured OT protocols. By storing more data locally, more advanced analytics can be pushed to the edge. Machine learning techniques and applications can be applied to predict outcomes like machine failures before they happen. In more challenging communication environments, or in environments rich with streaming data that require real-time alerts and responses, more analytics are likely to get pushed to the edge. In other environments, it might make more sense to simply reduce, consolidate, and encrypt the data at the edge and let the cloud do the majority of enhanced analytics.

Meet ThingWorx

Building a complete system for the industrial IoT can be a huge undertaking, even for the most resourceful companies. It requires a lot of expert know-how, time, and capital – and in the end, many companies end up plagued by long IT project cycles and low return on investment. The ThingWorx platform was purpose-built from the ground up to fundamentally reinvent how organizations connect, analyze, manage, and experience all the "things" in the connected world.

Designed for industrial IoT applications, the ThingWorx platform helps organizations to meet and exceed their goals, including:

- Investing more time leveraging the organization's understanding and expertise to innovate above the existing infrastructure, building applications that leverage the underlying data rather than needing to understand the nuances of operating environments, connecting existing disparate systems, competing standards, incompatible protocols, etc.
- Achieving "time to insight" as quickly and cost-effectively as possible regardless of scale, to ensure the IoT solution is capable of delivering that first value, and also be architected to quickly adapt to new requirements and challenges as solutions are added, deployed, and enhanced.
- Delivering the best product or service possible in highly competitive environments

ThingModel - What Makes ThingWorx Different

While there are a lot of products today that can solve the technical challenges inherent on the edge, most provide little to no assistance with the implementation. They need to be stitched together with whatever else exists in the environment. Given the depth and breadth of devices, interfaces and protocols in the OT, this can quickly become a daunting, resource intensive, and expensive challenge.

ThingWorx is built upon an enhanced entity-relationship model-driven architecture known as the ThingModel. The ThingModel is the digital representation of an actual physical "thing" – a true digital twin of a device or process made up of real-time data on that thing's properties, services, subscriptions, and events. It can be built from an unlimited number of data sources - structured and unstructured, time series, in motion and at rest. It employs pre-defined templates that can be propagated across like entities. It automatically generates an RESTful API, enabling it to be easily integrated throughout every ThingWorx module - ThingWorx Analytics, ThingWorx Utilities, ThingWorx Studio, ThingWorx Industrial Connectivity – as well as any other 3rd party application and technology.

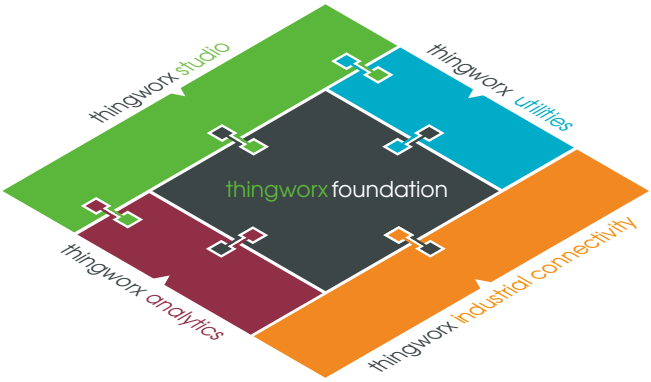

The power of the ThingModel is in its creation of an abstraction layer above the specifics of individual devices and protocols, based on binding device, data, and process relationships. This enables organizations to seamlessly tie in all of these components as they see them and employ them, substantially simplifying device and inter-system connectivity. The ThingModel connects and powers all other IoT platform modules and ensures the consistency and reliability of thing-related information by whomever uses it – both industrial IoT platform and application developers and the business experts. In turn, it rapidly breaks down many of the information siloes within the organization and creates a new framework for solving the challenges of integrating OT and IT by democratizing both the creation and enhancement of analytics and insights.

As a result, ThingModel enables the rapid creation of a broad range of IoT features and functionality rather than using manual coding or trying to connect disparate frameworks or pieces of technology. It is powerful enough to meet the needs of the most skilled developer, but it also powers the drag-and-drop functionality of a simple user interface that eliminates the need for coding and speeds the creation of high quality applications, dashboards, workspaces, and mobile interfaces. It enables

developers to easily integrate powerful capabilities – including analytics and augmented reality – into the solutions they create. And developers can confidently deliver a very rich repository of data that nontechnical business users can select, modify and enhance without concern for platform stability or performance.

Connectivity with ThingWorx

The ThingWorx platform is somewhat unique in its approach to providing device connection. Most IoT products and platforms today focus on the communications equipment connecting remote devices - like IoT Gateways. While this might be convenient in simplifying initial device connection, it forces all interactions on the OT edge to the gateways.

ThingWorx includes a comprehensive set of preconfigured templates and toolkits that address nearly any imaginable connection requirement. Virtually all ThingWorx connections support the unique AlwaysOn connectivity. Transport-agnostic, AlwaysOn utilizes a secure, persistent WebSocket interface to deliver encrypted, bidirectional access through firewalls for modeling, application tunneling, and audited file transfer capabilities.

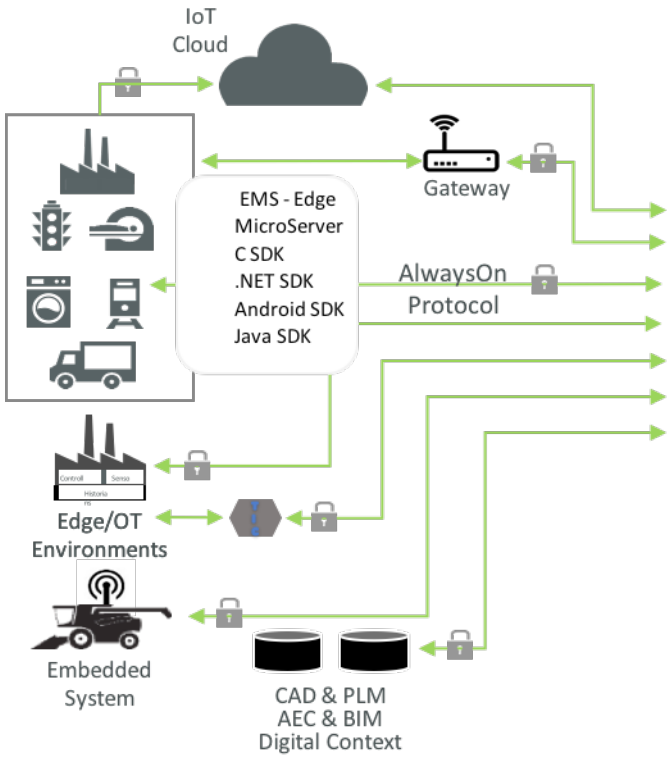

The **Edge SDK for ThingWorx** contains a complete set of libraries for building connectivity into a single standalone device or gateway. Organizations can choose from a set of preconfigured templates or easily create their own custom device agent using a number of languages including Java, C, .NET, iOS, and Android to integrate directly with the device application in its native language.

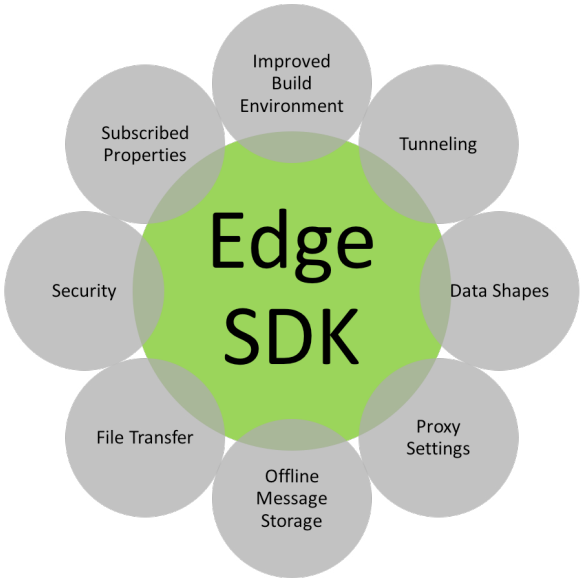

The **Edge MicroServer (EMS)** is a prebuilt, full-featured, lightweight IoT Gateway application that communicates directly with ThingWorx edge components whether they are running directly on the device, on a gateway talking with the device, or through a network interface. EMS enables bidirectional file transfer, remote software updates, and current actual state representation. It integrates with device applications on the edge through either a RESTful API, a Lua Script Resource, or a .NET Resource to all other components on the edge. There are

specific prepackaged edge components to enable simple integration to common industrial protocols like OPC-DA and OLE-DB to provide connectivity with HMI, SCADA, and Data Historians

ThingWorx Industrial Connectivity powered by Kepware, leverages OPC and IT-centric communication protocols to provide a single source of industrial automation data to all applications. This provides native communication and translation support for many of the systems and protocols commonly used in the OT to connect thousands of devices including machines, meters, PAC's, PLC's, RTU's, sensors, and more. Users can connect, manage, monitor, and control diverse automation devices and software applications across multiple transport options - including all radio and modem types, TCP/IP, Serial Multi-drop, and satellite - through one intuitive user interface. Streaming real-time data to SCADA and MES systems as well as other enterprise applications enables improved operations and decision-making. ThingWorx Industrial Connectivity streams real-time tag data directly into the ThingWorx platform.

The **Integration Framework** for ThingWorx substantially reduces both the time and complexity involved in connecting to other critical applications like PLM, ERP, and CRM. The framework automatically creates a ThingModel and integrates 3rd party data through their RESTful API into ThingWorx using standards-based integration connectors including Swagger and OData. ThingWorx Foundation converts all 3rd party data to standard ThingWorx format, transparently making this data available to all ThingWorx modules including ThingWorx Studio and ThingWorx Analytics. There are many integrations for OT applications already available including National Instruments, OSIsoft, CODESYS, Ansys, SAP ERP, HPE, Dell, and Cisco.

Connection Servers for ThingWorx provide deep interoperability with the major public cloud service providers including Amazon Web Services, Microsoft Azure and GE Predix. This integration provides the best of both worlds – support for the organization's existing cloud provider integrated with the rapid development-enabling features of the ThingWorx Platform. ThingWorx is ready to connect and offer scalable, secure, fully functional integration for the most remote and challenging devices, regardless of connection technology or provider.

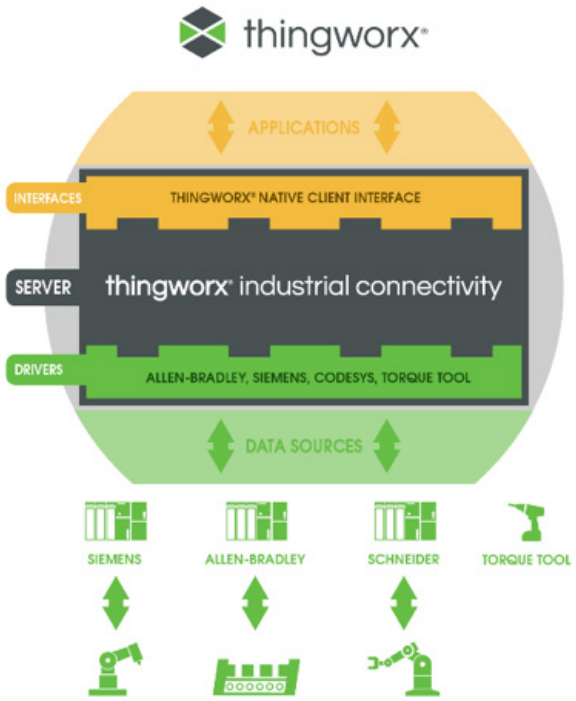

Deployment options for ThingWorx range from the smallest deployment on the Edge involving a handful of connected devices to the largest-scale customer deployments through a global scale architecture. ThingWorx can be deployed on-premise, in public or private clouds, or using a hybrid approach. It provides transparent connectivity, working either on or with public clouds, and is designed to provide continuous, secure bi-directional connectivity between devices and the ThingWorx server regardless of network architecture. And, with the growing prevalence of fog, dew, and edge computing, the ThingWorx simplifies the management plane with its federated and distributed architecture - ensuring software, firmware, data movement, and the security chain are all updated, linked, and synchronized.

The ThingWorx platform supports full federation, enabling sharing of the workload between ThingWorx servers. When the environment requires server clusters or regional servers, peer-to-peer connections can be established through federation to share and analyze data. ThingWorx-enabled things can share properties and services between servers. Federation simplifies distributed and tiered data storage and analysis. It provides autonomy at the edge and ensures the analytics and actions required in near-

real time are available while optimizing bandwidth. A central ThingWorx server can connect to each plant-level server, pull data, and aggregate it for regional or corporate level views. As users drill down into greater detail, the plant-level servers can propagate the required data to the central server.

ThingWorx employs a unique matrix multi-tenant design for data consumption and analysis. Most IoT platforms provide only limited multi-tenant visibility as they require mutually exclusive dataset definitions. ThingWorx Foundation will support mutually overlapping data sets, enabling enormous flexibility and granularity in defining user-based roles for mashups, reporting and augmented reality.

ThingWorx Analytics

ThingWorx Analytics delivers powerful, automated analytics capabilities including real-time pattern and anomaly detection, automated predictive analytics and contextualized recommendations. Utilizing simple user-friendly interfaces, visualizations, and easy-to-use tools, ThingWorx Analytics eliminates the need for developer or user expertise in data modeling, complex mathematics, statistical analysis, artificial intelligence, or machine learning. ThingWorx automates anomaly detection in real-time through machine learning and artificial intelligence. Anomaly detection can be set for any sensor or device in any location on the edge or in the cloud. Anomaly detection observes and learns the normal state pattern for each device without the need for setting rules or applying pre-calculations. When it detects any deviation, it delivers real-time alerts and provides a degree of certainty based on its cumulative learning.

Considerations for Moving Forward

It is important to remember that organizations today have a lot of flexibility and choice to build and optimize their industrial IoT infrastructure. The first priority should be to define a common set of the organization's goals - collaboratively developed by critical stakeholders in OT, IT, and other affected business disciplines - that addresses the specific requirements and the economics of the organization. This establishes the framework for moving forward, sets the priorities for where to start, highlights the skillsets needed to get there – both internally and externally through partners, and defines the benchmarks used to measure results.

Organizations should resist the temptation to simply experiment ad hoc to solve a single problem or provide a single solution. Many organizations have systems in place with new, integrated IoT capabilities that have been added by the manufacturer of that system. Without the benefit of a common set of organization goals, these systems could pose a looming integration challenge for many IT departments in the near future. Best case, everything works just fine. More likely, the organization must decide whether to forgo the capabilities, accept them as they are, or consider a costly transition to a different solution.

Developing a clear understanding of where the organization is today and a game plan for how to proceed will help to identify the best long term partners and solutions, both for the edge and throughout the entire organization.

© 2017, PTC Inc. (PTC). All rights reserved. Information described herein is furnished for informational use only, is subject to change without notice, and should not be taken as a guarantee, commitment, condition or offer by PTC. PTC, the PTC logo, Product & Service Advantage, Creo, Elements/Direct, Windchill, Mathcad and all other PTC product names and logos are trademarks or registered trademarks of PTC and/or its subsidiaries in the United States and other countries. All other product or company names are property of their respective owners. The timing of any product release, including any features or functionality, is subject to change at PTC's discretion.