CONNECTED INDUSTRY
GROWTH PARTNERS™

Advisory | Executive Search | Ventures

WHITE PAPER

The Intelligent Edge: Opportunity and IoT Innovation

By Ed Maguire
with assistance from Cate Lawrence, Paul Anderson, Jesse Demesa, Jim Fletcher, and Sandra Mueller

Executive Summary

The concepts around Edge computing are not new but are increasingly relevant as IoT-connected systems scale and become increasingly autonomous. Interest in, and demand for more powerful real-time data analytics fuels interest in an evolving architectural paradigm that harnesses the advantages of powerful centralized data centers with the nimbleness and flexibility of devices positioned at the physical edge.

There are technology and business reasons driving interest in variants of Edge computing. With often overwhelming volumes and velocity of sensor and endpoint data, it's neither cost effective nor practical to rely on centralized data centers for processing and analysis. Edge computing enables raw data to be filtered and aggregated as it's generated, with only the most important data sent to the data center for deeper analysis. This approach reduces **the costs of data ingestion and network latency.**

Edge computing makes sense for business reasons as well. When there's need for real time responses to sensor data, for instance on a connected vehicle, oil well or manufacturing machine, Edge computing enables data to be analyzed and acted on in near-real time. Edge computing is poised to enable a powerful new generation of flexible and distributed analytics for the Connected Industry, ultimately powering a broad range of autonomous and self-optimizing systems.

We wrote this paper to address the growing interest and curiosity around Edge computing. We endeavor to provide an overview of drivers, definitions, benefits and challenges of Edge computing technologies.

Contents

2. Executive Summary
4. Introduction
5. Surging Industry Interest in Edge Computing
6. The Industrial IoT Journey Is Still in Its Early Stages - Back to First Principles
7. The Coming Data Deluge
9. Back to the Future Again - Cycling from Centralized to Decentralized Paradigms
10. Defining the Edge
11. The Key Role of the Intelligent Edge
12. Beyond Cloud and Fog to “Edge Hives”
14. The Role of the Edge Device in Real-World Scenarios
14. One Size Does Not Fit All
16. Security Belongs Top of Mind
17. Looking Forward: The Promise of Blockchain/Distributed Ledgers
17. Edge Computing as the Future of Hybrid IoT Architecture
18. Use Case - Litmus Automation
20. Use Case - Sight Machine
21. Conclusion

Introduction

There's growing demand in the industrial world to benefit from compelling, substantial real-time insights 'under the hood' – from connected cars, factory machinery, agricultural processing, wind farms, oil and gas refineries and any other types of connected equipment. A sensor enabled machine on the factory floor produces data that provides production managers with previously unobtainable real time insights into the operation of the machine - predicting flaws in the machines before they occur. Precision analytics can reduce downtime, reduce operational costs and increase operational efficiency leading to substantial savings.

The more connected devices, more data gets collected. As a result, capturing, aggregating, and analyzing data quickly and efficiently becomes a greater challenge. With the emergence of autonomous vehicles, drones and data-intensive industrial automation systems, **the prevailing centralized cloud-mobile computing model is not designed to accommodate the extremely low latency demands of real-time, high-volume analytics.** As a result, there is growing interest in intelligent edge devices – gateways and devices that incorporate real-time data processing, advanced analytics and business logic.

Beyond latency, there are other technological demands: sheer volumes of data, the need to selectively deliver data to the cloud, benefits of solution flexibility, costs of networking and other factors. **While cloud computing remains foundational for Industrial IoT, a new cohort of Edge computing addresses when data needs to be processed locally in real-time.**

There are two distinct but complementary analytic roles: Intelligent Edge devices provide real-time analytics and decisioning, while cloud or on-premise data centers deliver in-depth insights into trends based on analysis of historical data.

Surging Industry Interest in Edge Computing

There's growing industry momentum around Edge computing. **In June 2016 Cisco and IBM announced a global collaboration to provide instant IoT insight at the edge of the network targeting use cases where time is of the essence,** but bandwidth is often lacking **such as oil rigs, factories, shipping companies and mines. In November 2016, the Edge Computing Consortium (ECC) was launched by** organizations including Huawei, Shenyang Institute of Automation (SIA) of Chinese Academy of Sciences, China Academy of Information and Communications Technology (CAICT), Intel, ARM and iSoftStone.

The Linux Foundation* in April 2017 announced EdgeX Foundry,** an open source project to build a common open framework for IoT Edge computing and ecosystem of inter-operable components **aligned around simplification and standardization** of industrial IoT Edge computing. **Later last year,** the ***Industrial Internet Consortium (IIC) established the Edge Computing TG to define Edge computing reference architecture. AT&T detailed plans for an Edge computing test zone in Silicon Valley to enable developers to trial use cases such as connected cars and VR.

ARC Advisory Group surveyed over 300 end users in 2017 from IIoT industries such as mining, machinery and automotive, and found that the majority of respondents (63%) were invested in the edge in the current business operations: Of the 63% of respondents 34% were either conducting a pilot; implementing a production level edge based system (17%) or getting edge enabled results (12%). **User interest was largely driven by operational interest in applying edge-based analytics. The primary reasons for consideration were to improve asset performance and maintenance (**considered by many to be the breakthrough application for the IIoT), particularly among process respondents (especially in North America and EMEA) and the drive to improve and optimize production, i.e. to prevent unplanned downtime.

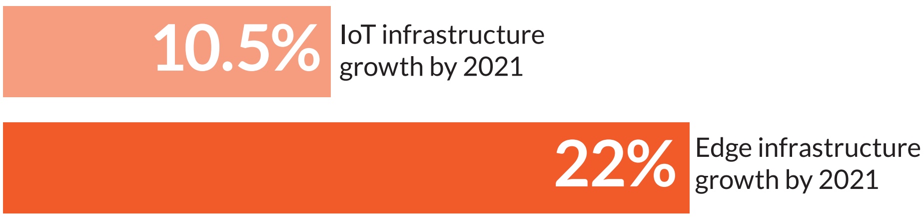

High growth opportunities will come from the Edge

There is promising growth forecast for Edge computing adoption. IDC estimates that by 2019, 40% of IoT Data will be processed at the edge, while the overall market for IoT infrastructure is forecast to grow at a 10.5% Cagr to $16 billion in 2021, Edge infrastructure is forecast to grow at a 22.0% Cagr to $3.4 billion in 2021. Growth in compute will primarily come from edge gateways and converged IT/OT systems.

High growth opportunities will come from the edge

by 2019, 40% of IoT Data will be processed at the Edge

Edge infrastructure is forecast to grow at a 22.0% Cagr to $3.4 billion in 2021

The Industrial IoT Journey is Still in its Early Stages

Back to first principles

The promise of IoT is the realization of incremental business value derived from applied analytics, powered by data collected from connected "things". Analytics create value in three fundamental ways: by managing risk, optimizing operations and growing the value of the solution. Successful implementations of IoT analytics can go much further; beyond improving specific aspects of a business, IoT analytics can drive fundamental business transformation.

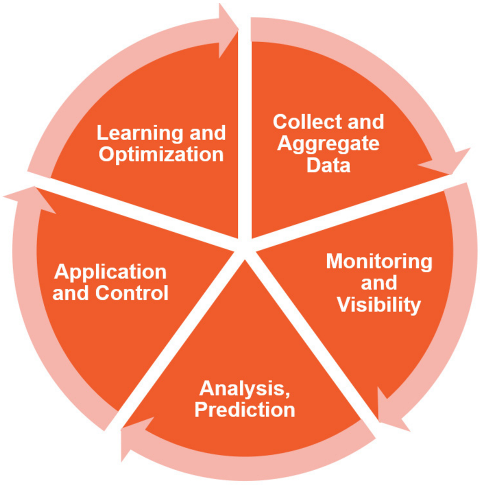

Figure 1: The Analytics Life Cycle; Source: Momenta

There's an inherent value cycle in analytics that's been consistent across the evolution of the market, stretching back to the rise of data warehousing. This involves turning data into insights into actions that drive business optimization. The initial step in the process for IoT analytics is instrumentation of physical assets and processes – putting sensors on things to generate data about what they do, and what their conditions are.

Taking a broader view of the opportunity around IoT adoption and maturity, **there's a series of phases as companies undertake their analytics journey. Each successive stage builds upon the prior phase.**

1. Initially the first connected asset use cases involve simple "connect and collect" providing basic visibility and monitoring – establishing a pulse on what's happening in the physical world.
2. By leveraging the wealth of data collected, proactive analytics help identify conditions or potential failures that need to be addressed.
3. Advancing deeper into the analytics world, predictive analytics attempt to pinpoint pending failures or problems, so they may be addressed in advance. Prescriptive analytics actually guide specific steps necessary to address them.
4. The most advanced phase is autonomous systems that can self-diagnose and guide actions to keep systems running seamlessly.

As we look at the introduction of Edge computing, elements of the data management and analytics can be performed within the cloud, at the edge, or somewhere in between.

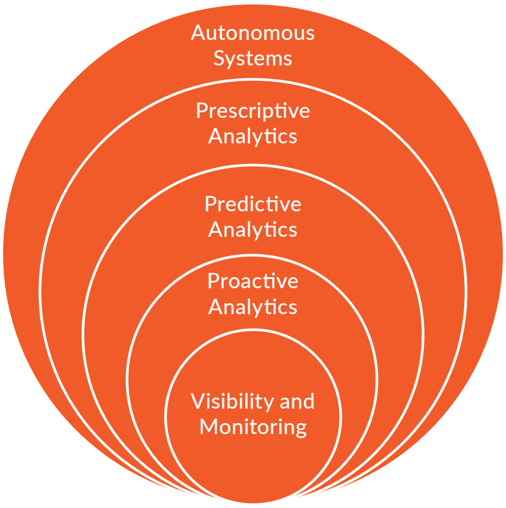

Figure 2: The Evolution of Analytic IoT Systems; Source: Momenta

As visions mature into reality, the advance of more sophisticated analytic solutions including predictive and prescriptive analytics and ultimately autonomous systems will drive need for lower latency, more powerful analytics. Intelligent Edge computing will empower a future of autonomous systems.

The Coming Data Deluge

The next decade will see an acceleration of data creation with the connection of more devices, and more types of data. **The combination of nearly universal connectivity, low cost storage, and low-cost compute will change the game forever.** Over time, complex industrial machines and vehicles will effectively become mobile data centers.

Data gets bigger and bigger

	PER DAY
Connected Car:	4 TB / per car
Oil Rig:	10 TB / per well
Solar PV Systems:	5 PB
Jet Engine:	2.5 TB
Wind Turbine:	5 TB
Ship:	1.5 TB

As IoT connections scale, more industries get serious about adoption.
Steadily and relentlessly, there will continue to be more and more devices connected to networks. IDC predicts **the worldwide installed base of IoT endpoints will grow from 14.9 billion in 2016 to over 82 billion in 2025.** Gartner estimates 8.4 billion devices were connected in 2017. After a few years of relatively modest adoption in industry, the pace of industrial IoT connections is accelerating, according to Verizon's 2017 State of the Market IoT Report.

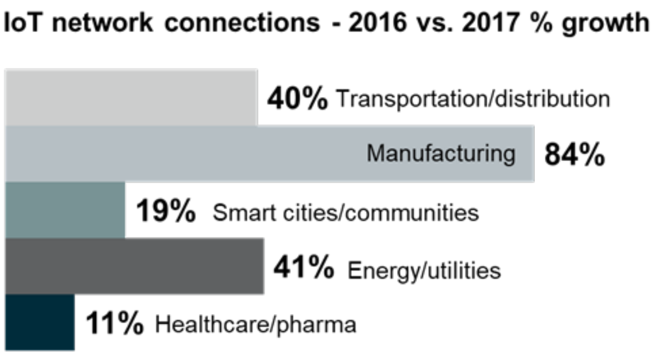

Figure 3: Year-on-year growth in Verizon IoT connections

Data Overwhelms the Centralized, Cloud-based Model

With the massive influx of data, there are several current challenges relying solely on a mobile-cloud-based approach for IoT. A cloud architecture where data, logic and processing reside in centralized data centers creates several choke points for IoT data, particularly in increasingly dynamic scenarios.

There are limitations to relying solely on a centralized cloud-based approach for IoT:

- **Data latency:** In time sensitive use cases which require real time data analytics, such as autonomous vehicles, noxious gas monitoring, safety equipment, and other scenarios where a split-second delay of data going to the cloud and back to the device could be disastrous or deadly, analytics need to be embedded into operational processes.

- **Locational challenges:** Localized processing needs to take place in diverse environments - hardware needs to be rugged to survive diverse environments including hard to reach locations, 'canary in mine' spaces and those exposed to harsh weather. This could include an offshore oil rig, underground mine or deep water well. These environments can be isolated resulting in unstable links with limited bandwidth and variable latency.

- **Customers don't want siloed systems or services:** In many industrial workplaces, the environment may be a digital divide; a cornucopia of unconnected legacy machines, non-compatible devices and controllers. Information exists in siloed systems and there's a physical in-house divide between IT and OT. Part of the reason is because we are in the early stages of a change of capabilities – the former generations of devices and protocols were aligned with the physical capabilities and environments. These technology limitations are rapidly disappearing, and the potential cost points for new technologies are significantly lower than the legacy approaches. However the challenge is that companies that have dominated traditional operational technologies have little incentive to see the market transform to a lower cost, and into different tech environments.

Back to The Future Again

Cycling from Centralized to Decentralized Paradigms

The principles of computing at the Edge are not new. M2M, SCADA, and on- premise industrial control systems have been established for decades, and so have concepts around data collection, analytics, and orchestration. **The primary difference is in the set of design assumptions, and the methods of application using the Intelligent Edge.** The state of the art of distributed computing, communications, storage, and analytics has changed significantly over time, allowing the application of state of the art IT concepts to be applied within an operational technology context, including on the Intelligent Edge.

Evolving Computing Paradigms

From Centralized to Decentralized - Usage scales by orders of magnitude

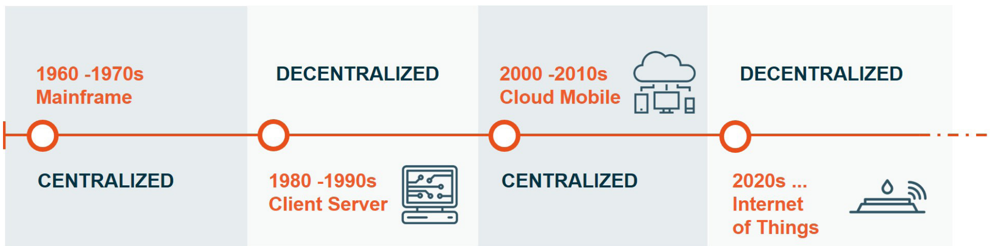

Figure 4: Evolving Computing Paradigms

Over time, computing eras have alternated between centralized and decentralized, and with each paradigm shift, usage increases by orders of magnitude. The mainframe/minicomputer eras were about centralized computing with tens of thousands of users. The PC/client-server era became decentralized as the number of users increased into the hundreds of millions. With the rise of today's cloud-mobile world is largely centralized. Today, Cloud is where much processing gets done, with mobile devices providing the front-end interface, with the number of users scaling into the billions. With IoT adoption, scale of devices will be reaching into the tens of billions and beyond and this is driving the shift to a decentralized paradigm.

Defining the Edge

Edge computing is a broad term that straddles several broad categories, and definitions can be fluid, so we'd like to put a finer point on it. The IoT data flow begins with front-end sensor data and ends with systems using analytics and decision-making capability driving intelligent actions. We look at Edge Computing technologies along a continuum:

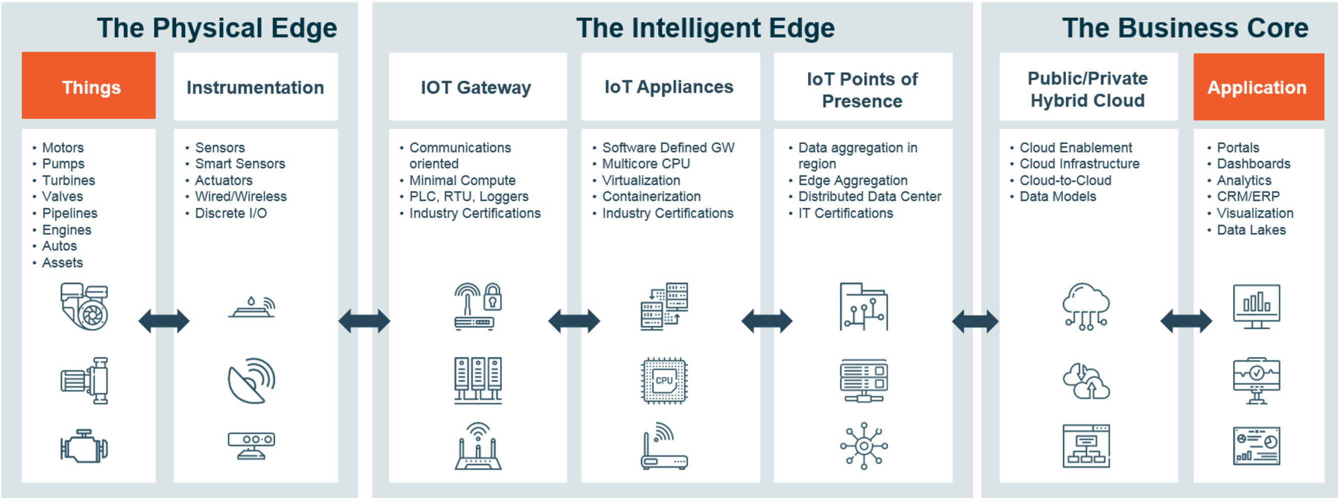

The Physical Edge: This can also be referred to as the "hard" edge, the Physical Edge includes connected "things" as well as devices, sensors, actuators and other machines and devices that generate data from processes, instrument readings and other sources.

The Intelligent Edge: This comprises what most people think of when they hear the term "Edge Computing". **In our definition, the Intelligent Edge encompasses three variants of increasing size and scope.**

- **IoT Gateways** are simple edge devices, focused on communications functions, with minimal compute capabilities.

- **IoT Appliances** are far more robust, multi-core, often running virtualized or containerized applications, sometimes carrying certifications for specific industry use cases. IoT Appliances include communications and networking capability and also provide custom application and compute environments employing virtualization, containerization, and granular application management.

- **IoT Points of Presence** are super–aggregation points for other intelligent gateways, also referred to as Distributed or Edge Data Centers or Data Centers at the Edge. They represent another order of magnitude greater in scale. IoT Points of Presence can include multiple machines, extensive storage and processing capabilities with the goal of providing more robust processing and storage with low latency to and from physical edge.

The Business Core: This comprises the core of the enterprise compute infrastructure, whether it reside in a public, private or hybrid cloud, with all of the supporting data management infrastructure, compute and storage capacity to support business logic and analytics for the end users.

The Key Role of The Intelligent Edge

Analytic Must Produce Insights to Add Value

- Complex Event Processing (CEP) can be moved to the Edge
- Domain-specific analytics models can be developed, iterated, and quickly deployed
- Analytics will be used for the specific business domain and for managing the infrastructure
- Analytics must move from “words and numbers” to ”automated actions”
- Technologies such as Augmented Reality will become the norm for human visualization of complex information

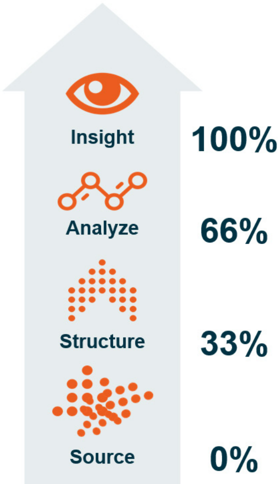

Middle tier gateways play an integral role in data management. These devices are typically data aggregators, with communications and networking capabilities. Middle-tier gateways are suited to manage high volumes of continuous data, returning only critical, business relevant data.

The challenges in how to best manage the movement of data from the front-end to the backend involve a range of choices around application architecture, communications protocols, storage, processing power and other considerations. There is no “one size fits all” approach for gateways, appliances and Points of Presence.

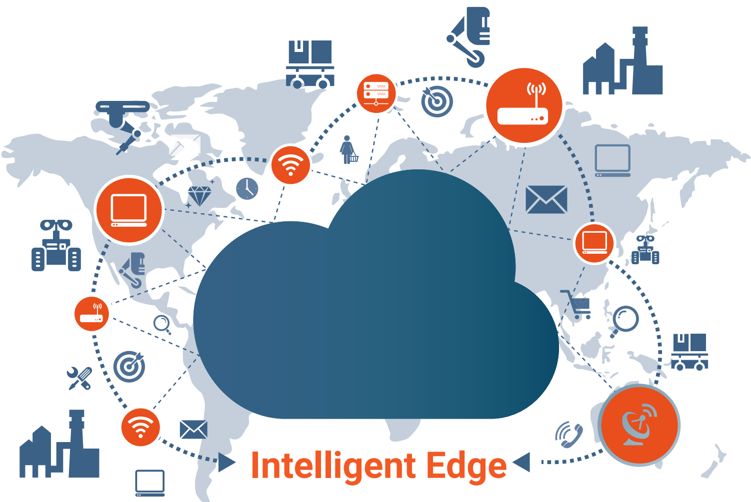

Edge Computing opens up new possibilities and enables scenarios that were not possible or effective under a centralized cloud approach.

A few advantages deploying edge-based devices include:

- **Speed** - Edge computing reduces latency because data doesn’t have to travel over a network to a remote data center or the cloud for processing, followed by another delay as actions are driven from the cloud.

- **Network independence and its locational benefits** – Edge computing enables the ability to operate independently of the network availability or cost, to drastically reduce data transfer rates, and the ability to control what data stays on-site vs what is delivered to the cloud.

- **Scalability and cost reduction** - Edge computing is fundamentally 'distributed computing,' meaning it improves resiliency and reduces network load, and can be more efficient overall when processing data at the point of origin. These characteristics all lower overall cost. Data transmission costs are lower because the amount of data transferred back to a central location for storage is reduced. Edge computing enables digital solutions with greater autonomy, reliability and security.

- **Security** – This can be a pro or con but needs to be mentioned. If the data remains at the edge, the edge has to be secured as well. Physical security becomes more key as data is retained on the edge.

A good use case example for Edge computing is the transmission of wind farm data: Wind farms can consist of hundreds of wind-powered turbines generating vast amounts of data. According to research firm Wikibon, a typical wind farm is embedded with security cameras and other sensors – and there is a potential physical distance of hundreds of miles between the wind farm and the cloud data center. Through processing real-time data at the edge, and transmitting summary data to the cloud data centers, one wind farm reduced traffic flow by 95% and reduced the cost of management and processing by 64% over three years.

Through processing real-time data at the edge, and transmitting summary data to the cloud data centers, one wind farm reduced traffic flow by 95% and reduced the cost of management and processing by 64% over three years.

Beyond Cloud and Fog To "Edge Hives"

While it may be tempting to champion Edge and fog computing as the panacea of all the pain points of cloud computing, the reality is that we are going to see more of an aggregate model, encompassing the isolation of the edge with the power of selective aggregation of data and "nodes" at the edge.

We refer to these models as Edge Hives - aggregations of processes that can operate autonomously and also can share/learn with other hives and clouds. These can be physical or logical aggregations, with software containers selectively and dynamically aggregated into hives. The hives can provide a model around which one can build out solutions at the edge - so that there's a succession, and each logical "hive" can have a different set of roles. Container technologies such as Docker become really important, because we can build out solutions in the same way as they were built out in the cloud, using a micro services approach to deploy them to the edge that provides a common development and deployment model across whatever the solution. The hive approach would allow us to dynamically repurpose or adapt the capabilities of the edge nodes as well.

For example, in the use of Edge computing in Oil and Gas, many companies also rely on drones, whose data is processed on cloud servers. According to Schneider Electric, drones are used in Oil and Gas exploration to locate promising sub-surface pockets that contain hydrocarbons. More specifically, drones are effectively flying sensor platforms, dispatched to conduct gravity, magnetic, topological, mapping and other surveys. Drones also are suited for surveillance or inspection of remote facilities, often performing their tasks at far lower cost than alternative methods. The data they gather is fed into the network and then into computing platforms. This information is digested and becomes part of the knowledge that determines the most promising sites to explore and pictures of areas that were previously inaccessible (or accessible only at a very high cost).

While many IoT devices require real-time decision making at the edge, businesses require historical analysis for process improvement and model development. This is best done when the data from multiple edge devices can be combined centrally to provide the relevant holistic view.

Moreover, insights gained from historical analysis can be pushed back to the edge, so that an IoT-enabled edge device continually evolves to make better real-time decisions. The compute model becomes a combination of Edge and cloud computing where IoT devices operate at the edge in real time, collect and process raw data at the edge, and share metadata to the cloud for comprehensive historical analysis and continuous process improvement.

The Role of The Edge Device in Real-World Scenarios

Over time, we expect that the edge gateway and appliance functionality will become integrated and/or embedded in a variety of industrial systems including controllers such as SCADA systems, automated manufacturing machines, smart power grids, wind turbines, magnetic resonance imaging (MR) scanners, undersea blowout preventers, smart streetlights and traffic lights; motor vehicles (cars and trucks), rail, shipping, aerospace, smart buildings and many other scenarios. Industrial firms have tremendous opportunities to derive efficiency benefits because of the ability to scale.

In an increasingly distributed architecture model, enterprises will move computing and data centers closer to the connected devices at the edge. In turn, enterprises will increasingly rely on DNS technologies that include intelligent traffic management to direct workload across such highly distributed edge architectures. Furthermore, traditional industries, like manufacturing, that have been struggling to showcase differentiated products, will now embrace edge analytics to drive new revenue streams and/or significant yield improvements for their customers.

One Size Does Not Fit All

There are multiple ways to deploy edge technologies in real-world scenarios, and there's no single approach that fits every type of situation. The below taxonomy provides a helpful framework to think about the different elements involved.

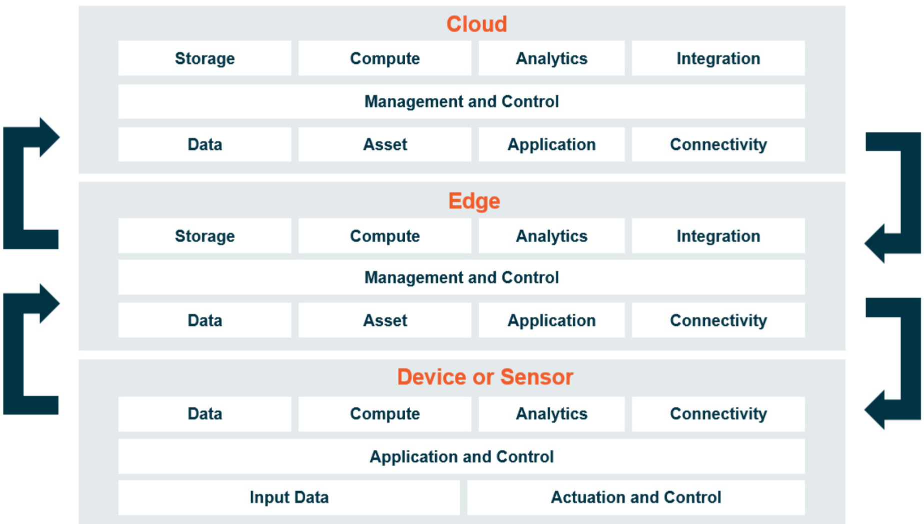

Source: Paul Anderson

We can see illustrating different types of scenarios there are different levels of sophistication. In the case of a simple environmental monitoring solution, data flows one way, from the sensors to the cloud via a low power network, and it's not critical to maintain persistent data connections from the end points to the cloud.

Simple: Low Power Secondary Sensing

Environmental Monitoring

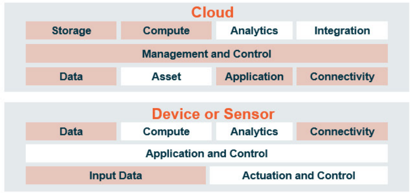

- Low Power
- Unidirectional Data Flow
- Low Data Rate
- Not Mission Critical

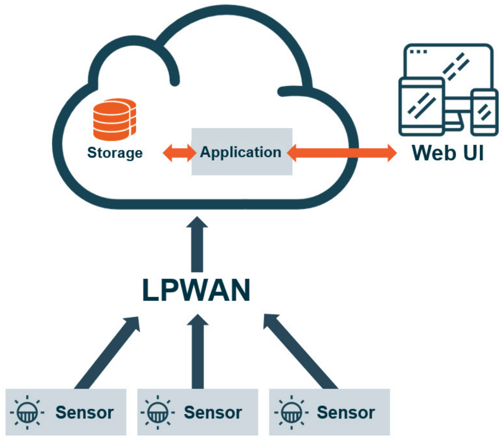

Source: Paul Anderson

A more complex scenario involves a remote site monitoring solution, such as a municipal well pump or lift station. In this scenario, data is fed to collector/ controller edge devices which incorporate limited capabilities for offline operations to be able to handle extreme situations.

Intermediate: Remote Site Monitoring and Control

Municipal Well Pump or Lift Station

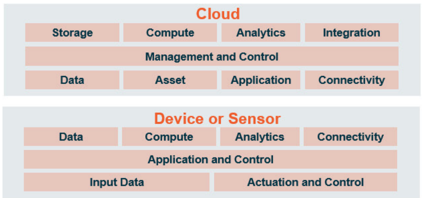

- Industrial Environment
- Medium Data Rate
- Limited Offline Operation
- Some Data Reduction and Local Analytics Possible
- “Embedded Edge”

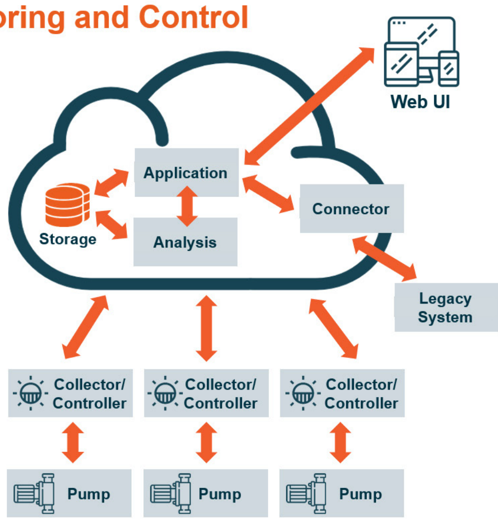

Source: Paul Anderson

In more sophisticated scenarios such as factory floors, there are many more technology challenges involved. IT and operational technologies need to be effectively integrated, edge data needs high-rate connectivity, and there needs to be full offline operational capabilities. Edge deployments are sophisticated scenarios in manufacturing or industrial settings that can incorporate multiple tiers of data collection, aggregation and processing – along with sophisticated orchestration of logic and redundancy to ensure that mission critical operations benefit from real-time insights while being cushioned from the impact of hardware or connection failures.

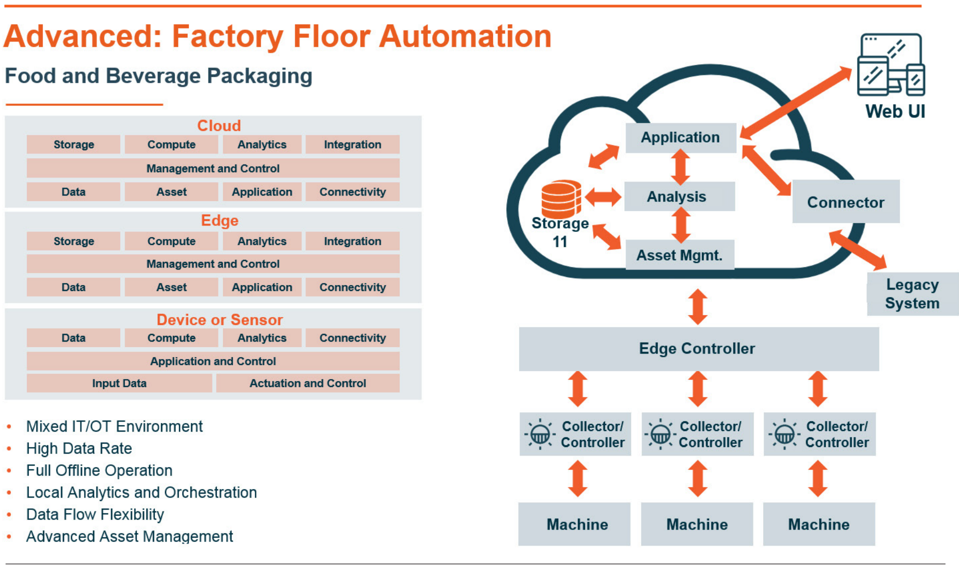

Source: Paul Anderson

Security Belongs Top of Mind

Security needs to be designed in from the start. This includes physical security, transport layer security, data encryption, key management, access rights and controls, management and of course threat prevention.

Traditional industrial SCADA systems enjoyed “security by obscurity” because machines weren’t connected into the Internet or the data center, or if they were, there were multiple levels of interaction to get there. Security for Edge computing involves several key considerations.

On one hand, data on an edge device does not travel over a network where it’s prone to interception. Data stored in enterprise or cloud datacenters benefits from established physical and virtual access controls, mature perimeter security defenses and security procedures.

Edge devices are subject to their own sets of vulnerabilities. Every connected sensor and actuator represents a potential point of compromise for a malware or DDoS attack. The Mirai Botnet attack in late 2016 commandeered hundreds of thousands of IoT devices in a large scale Distributed Denial of Service attack that brought down many top sites on the Internet.

Data encryption, access controls and use of virtual private network tunneling play important roles protecting Edge computing systems. Application logic running on edge devices also needs to be secured. Lastly, edge device authentication and identity need to be.

Many common IT security best practices can be applied directly in edge scenarios, while edge gateways and appliances can perform security monitoring and control functions to protect other devices and cloud applications.

An effective security strategy requires a holistic approach and consistent application of technology, configuration, and operations. Perimeter security alone is not sufficient given the complexity of IIoT solutions, with edge included or not. Each layer requires hardening, and solutions need to be evaluated end to end.

Data encryption, access controls and use of virtual private network tunneling play important roles protecting Edge computing systems.

Looking Forward: The Promise of Blockchain/Distributed Ledgers

As the volume of IoT connected “things” scales, there are growing challenges not well served by centralized architectural models. Blockchain technologies can serve to provide identity validation for things and devices, track transactions and manage payments between machines. Public, or permissionless blockchains provide economic incentives to operate network nodes that participate in validating, securing and recording transactions. The technology is still maturing, but there’s a lot of promise for use of blockchain and other distributed consensus technologies such as the Directed Acyclic Graph (DAG) used by IOTA, Hedera Hashgraph and others.

Edge devices such as IoT Gateways and Appliances and Points of Presence could be deployed as blockchain or DAG nodes. Because the entire record of a blockchain is stored on every node, encoding of identity or other security data could be accessed locally to accelerate response time for transactions such as micropayments or secure data exchange. Blockchain technologies can enable data providence for supply chain, privacy and other security functions, as well as smart contracts and micropayments.

Edge Computing as The Future of Hybrid IoT Architecture

There are financial, technological and business reasons behind the growing interest in Edge Computing. Edge Computing can optimize how to process data, in the most cost-effective way, with the most appropriate performance, for maximum business benefit.

There’s a highly complementary relationship between edge and cloud. For high value, time-sensitive data, edge processing at point of origin reduces the volume of data that needs to be moved, reduces network traffic and shortens the distance data must travel.

- **Edge computing reduces demands and dependency on the business core’s computing environment.**
- **Cloud servers and data centers can** process less time-sensitive data for historical analysis, reporting and data mining and can **store longer term data.**

Using the edge for real time data and cloud for historical analytics reduces costs and optimizes latency while improving the efficacy of mission critical applications.

Using the edge for real time data and cloud for historical analytics reduces costs and optimizes latency while improving the efficacy of mission critical applications.

There is a broad array of Edge computing options across gateways, appliances and points of presence. While there is no single approach that fits all scenarios, the business value of data driven decisions, and the speed with which they need to be made will inform the appropriate design for edge-cloud solutions. The technology is evolving, increasingly powerful and sophisticated – and we see the business impact in the coming era of autonomous systems as profound.

Use Case - Litmus Automation

Addresses Beverage & Brewing Company Challenges

Litmus Automation provides an extensive Platform as a Service for companies that are in a rush to embrace the disruptive Internet of Things technology and leverage it for real business challenges. Litmus simplifies the complexity of developing IoT systems and solutions with a secure and scalable middleware cloud platform called Loop. With Loop, companies can securely connect any type of hardware, device, sensor or machine to the Internet and integrate the data being collected to any 3rd party software application or enterprise system (ERP, CRM, Big Data Analytics, Dashboards, etc.) in real-time. Loop also contains an extensive device management suite for deploying and monitoring IoT systems. Litmus Automation focuses on the Automotive and Industrial industries and counts some large Fortune 500 companies as clients.

Background

From the manufacturing process of their product, to packaging and shipping, a multinational beverage and brewing company has a high demand for their products. Any malfunctions, resulting in delays to their to-market system, quickly become a costly situation.

Problem

Production line anomalies caused by old machine systems often require long-term and costly maintenance. It was common that the brewing process would be down for several hours when malfunctions would occur in their old machines.

Solution

Litmus Automation's Loop and LoopEdge platforms enabled predictive maintenance as a precautionary solution, eliminating the instantaneous cost of failure on old machine systems. Before a solution was implemented, Litmus Automation identified the single point of failure in the whole plant system with the client, which were the damper and damper driveshaft.

Litmus Automation used a Historian database interface to collect the data, as it was the client's only method of extracting data from its machines. LoopEdge pulled data from the Historian database interface using a time-series based algorithm that also standardized the data. The data was then pushed to the Loop platform using standard web protocols.

A machine learning algorithm ran on the top of the data to discover anomalies that indicated failures in the system. LoopInsights visualized the results of the data from the machine learning algorithm and anomaly detection. Normalized and processed data was also provided back to the client's database for them to conduct further custom analysis. In conclusion, the Loop platform enabled a predictive maintenance system that would alert the plant Floor Manager up to nine hours before an anomaly would occur in the system.

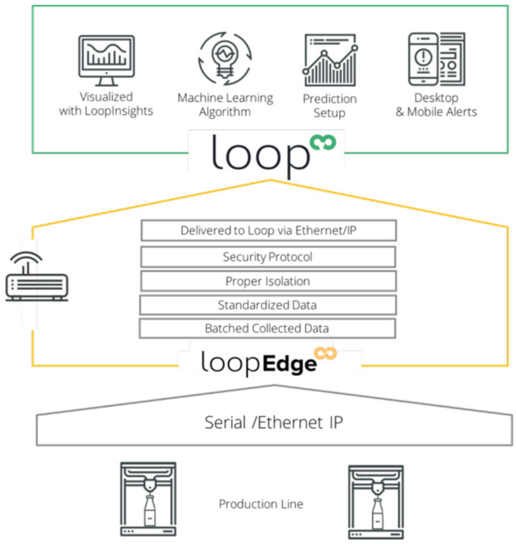

Source: Litmus Automation

Solution Journey

LoopEdge was installed at the Gateway level via a Historian system and batched data was collected in real-time. The data was standardized and pushed to Litmus Automation's Loop platform. A machine learning algorithm and anomaly detection system ran on top of the Loop platform, and LoopInsights displayed the data and results. Alerts were generated to notify the plant Floor Manager of any upcoming anomalies, up to nine hours in advance.

End Benefit to Company

The company was able to reduce and eliminate the instantaneous cost of failure using Loop and LoopEdge, saving time and money in the production process and optimizing maintenance schedules on their older machines.

Use Case - Sight Machine

Sight Machine

Sight Machine is powering the digital transformation of manufacturing. Used by Global 500 companies to make better, faster decisions about their operations, Sight Machine is the industry's only digital platform purpose-built for discrete and process manufacturing. With artificial intelligence, machine learning and advanced analytics, Sight Machine enables manufacturing leaders to answer the known and unknown operational questions of today and tomorrow.

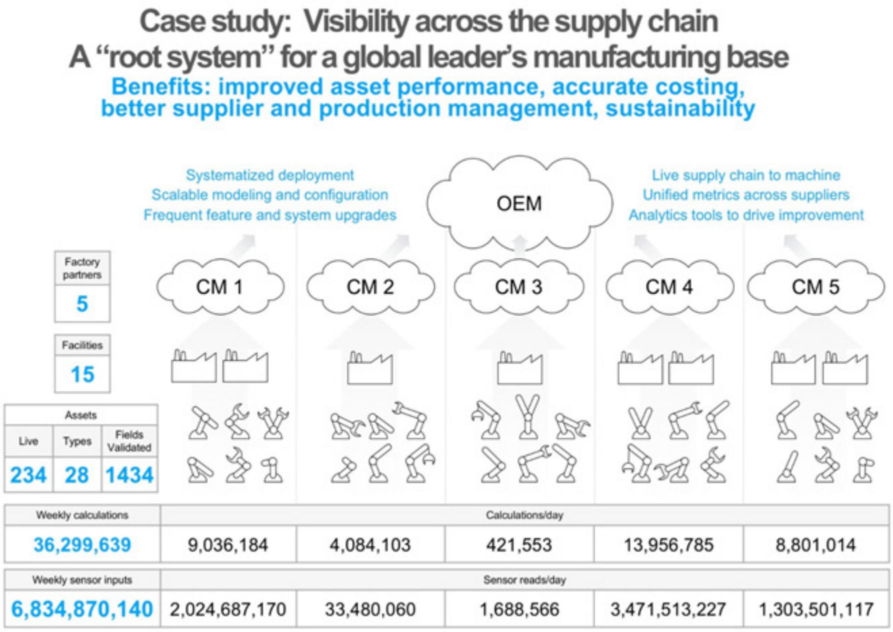

Source: Sight Machine

What True Transformational Capabilities Look Like

Problem

$32B manufacturer with 500 contract manufacturing facilities attempting to gain performance visibility across a wide network of plants, including contract manufacturers to better understand capacity constraints and climate the need to build/source additional manufacturing capability.

Solution

Sight Machine implemented across plants enables previously unattainable visibility and improved production consistency across the supply chain.

Benefits

Improved asset performance, accurate costing, better supplier and production management, sustainability.

Edge Computing is an important component to realizing the full potential of Industrial IoT. We've covered the drivers, definitions, and capabilities of Edge computing in this paper, endeavoring to convey appreciation of the complex dynamics involved with next-generation IoT architecture. Edge and cloud computing offer tremendous potential across a range of scenarios, and challenge for organizations will be to design and architect systems that not only harness the best capabilities available today, but also position themselves for future growth. Critical to success with Edge computing, as with emerging technologies, is to engage with trusted partners to chart the journey into the future.

Founded by deep practitioners in the Connected Industry space, **Momenta Partners** deploy their industry-leading Advisory, Executive Search and Venture Capital practices to accelerate the growth of Connected Industry companies and companies preparing for digital transformation.

3 Practices Hyper Focused on the Enterprise IoT

Momenta Advisory helps companies at all life stages capture the Connected Industry opportunity, providing strategic and operational guidance for digital transformation.

Momenta Venture invests in early-stage Connected Industry companies and advises clients in buying or selling strategic Connected Industry assets. Interacting with hundreds of Connected Industry companies worldwide, we have a pipeline of M&A targets and market knowledge inaccessible to our competitors.

Momenta Executive Search is the top retained executive search firm for Connected Industries. We have placed exceptional, strategic talent across all functions, within all major vertical markets in every continent for Fortune 500 leaders and venture-backed innovators.

CONTACT INFORMATION:

Ed Maguire
ed.maguire@momenta.partners

Ken Forster
ken.forster@momenta.partners

MOMENTA PARTNERS

Connected Industry Growth Partners™

Europe: +41 43 430 7516
New York: +1 (917) 765 3600
San Francisco: +1 (415) 366 6632

info@momenta.partners