Industry in Motion: McRock Perspectives on Industrial IoT

August 2019

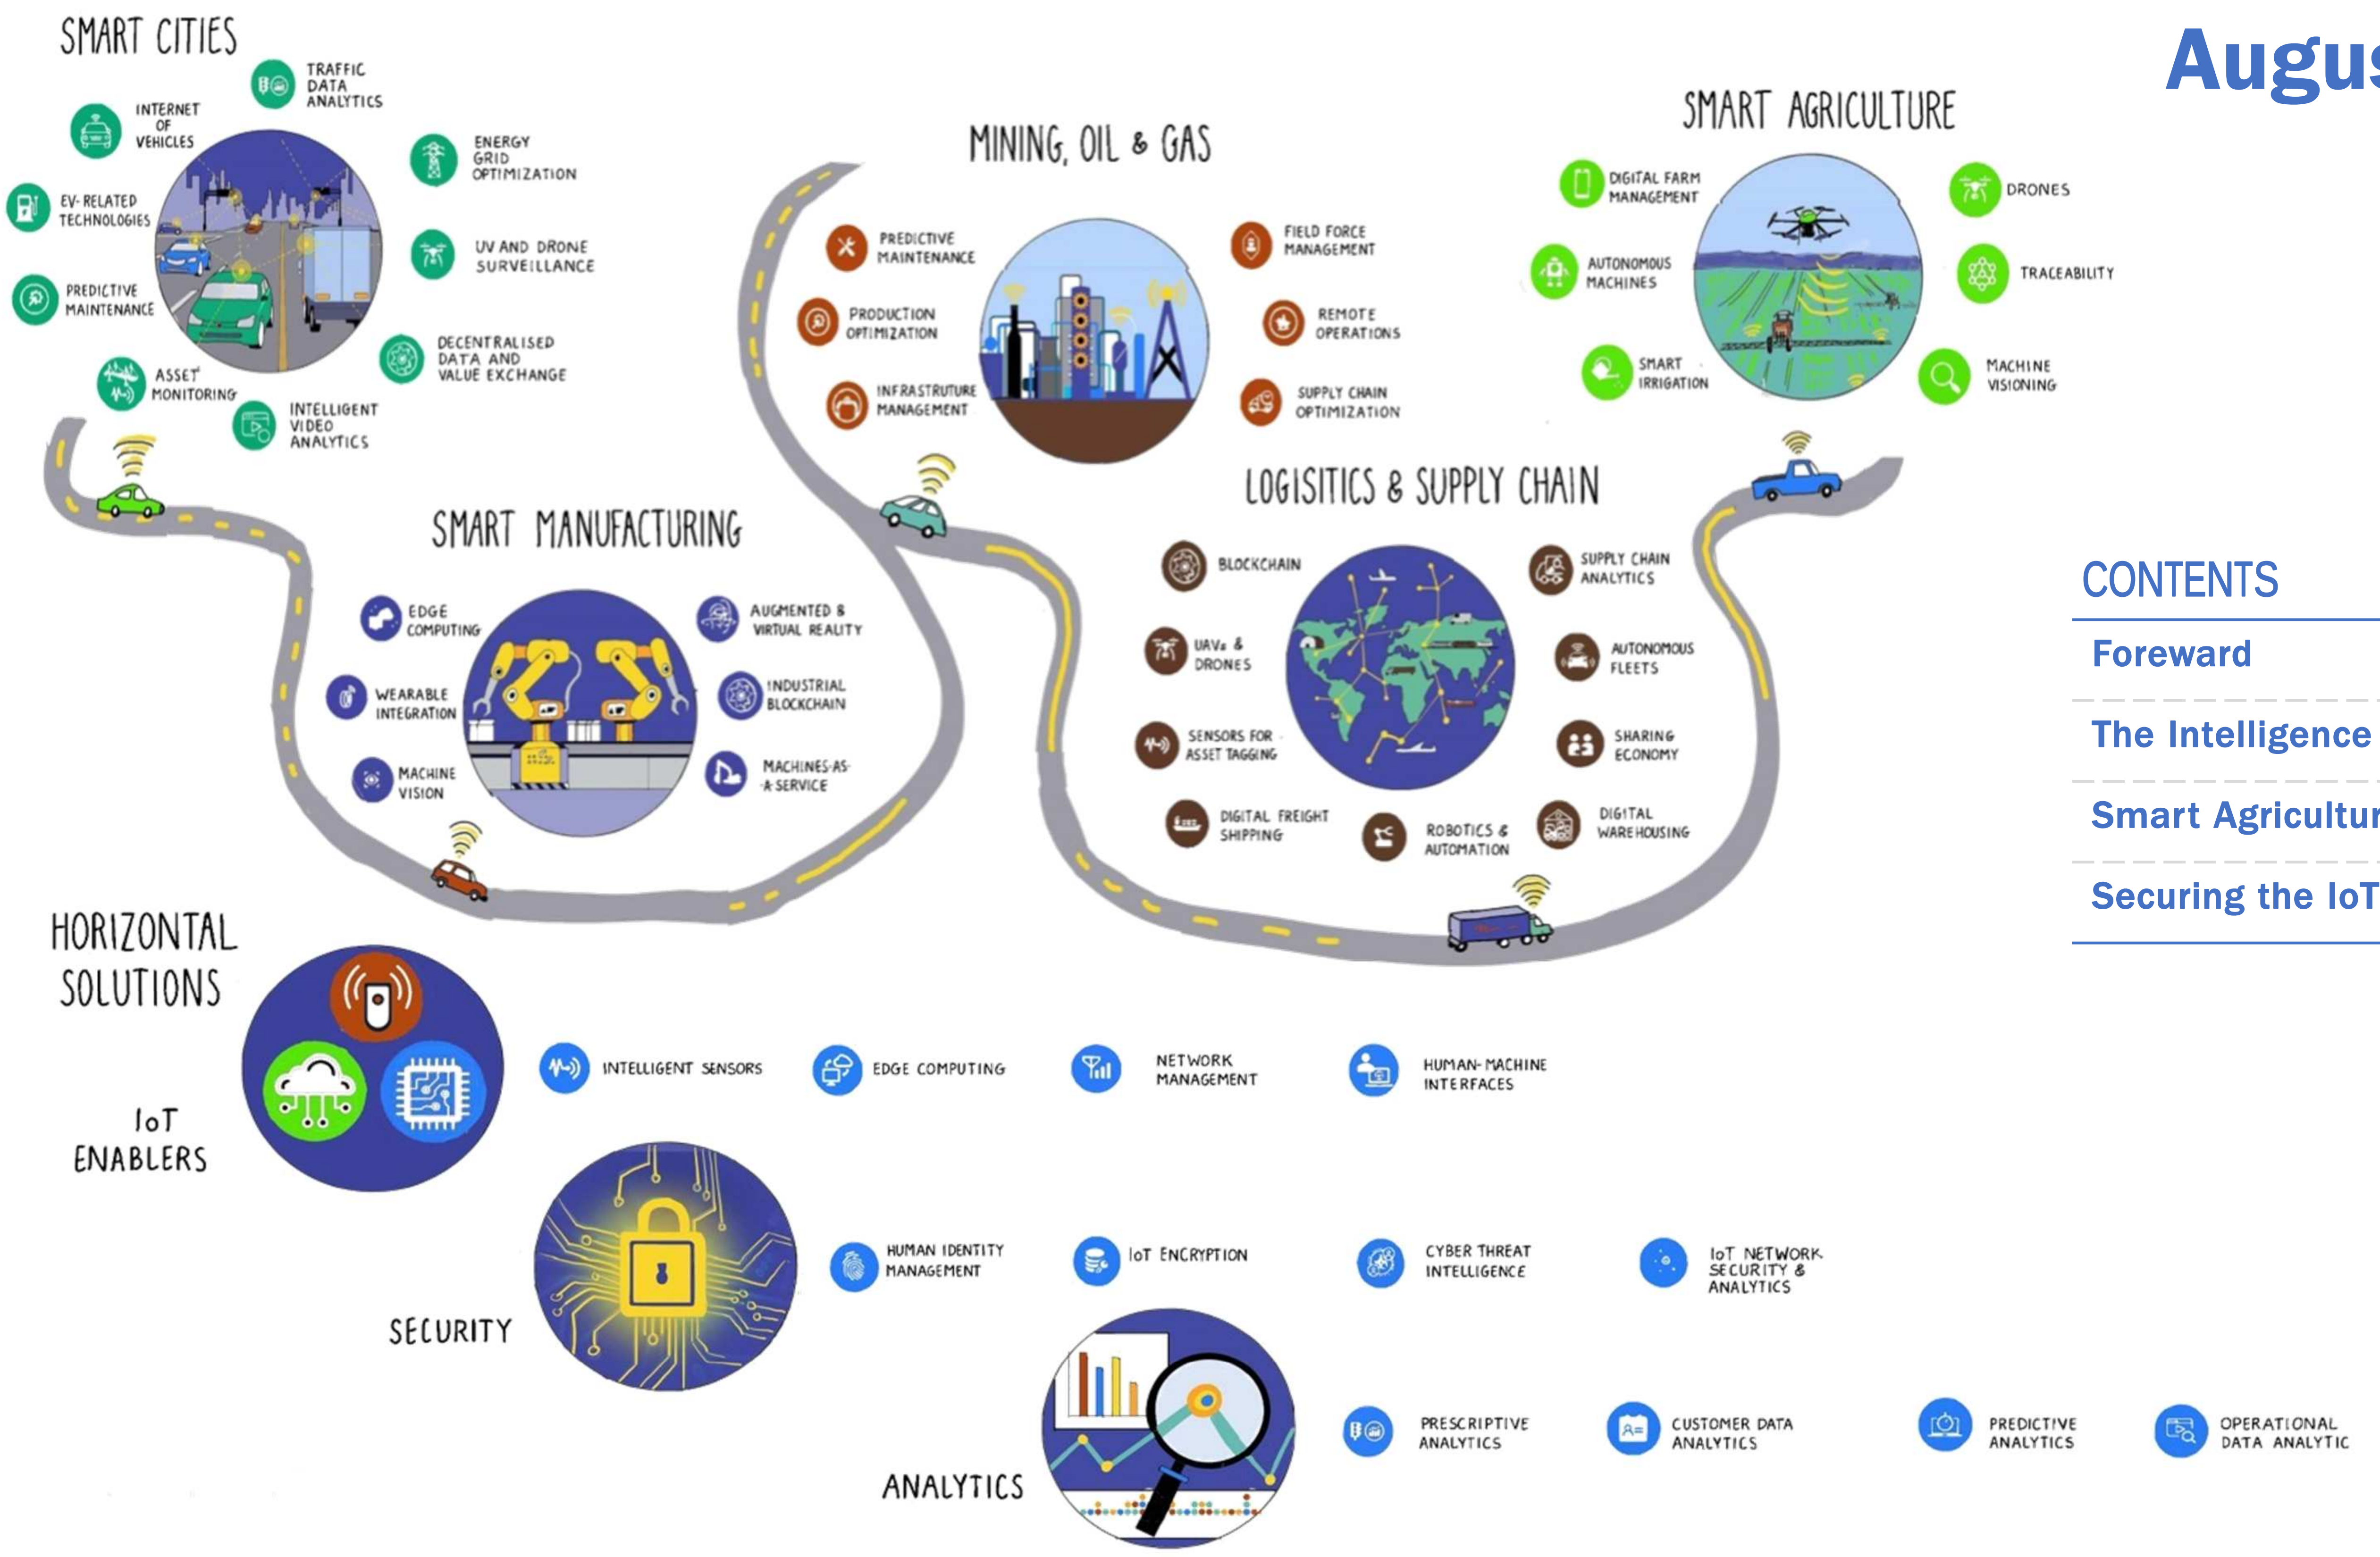

CONTENTS

Paula Hansen

McRock Capital

Foreward

The "Industrial" Internet of Very Important Things

As adoption of the industrial internet evolves, the transformation of data into enhanced business outcomes is increasingly interlayered by vertical idiosyncrasies and emerging horizontal forces. These complexities increasingly enter our focus as our attention shifts away from sensors ("building the body") i.e. creating access to volumes of data to successfully controlling and predicting complex business outcomes through Artificial Intelligence (AI) and analytics ("building the brain").

From the perspective of business process transformation, the use cases for IoT have never been clearer. We increasingly see incremental (leading to transformational) impact in areas such as predictive maintenance in manufacturing, where IoT technology can determine when machines will fail, or in smart cities, where there are obvious benefits to improving traffic flows and congestion.

As innovators seek to implement and scale these solutions, the success of Industrial IoT is increasingly determined by factors such as creative monetization - the inclusion of the broader ecosystem of participants (client-partners, OEMs, 3rd party developers) in tandem with new data-driven business models and revenue streams.

Industrial IoT: Value Creation

III. Lowering the financial risks of innovation for clients

- **Inclusive business models and revenue streams (e.g. As-a-Service)**
- **De-risking data**

II. Demonstrating value to clients

- **Increasing efficiency and cost savings for clients**
- **Partnering with investors**

I. Incubation

- **Sensor led technology**
- **Creating volumes of data**
- **Pilot projects**

Equally important is the emergence of a more secure data management layer, driven by data privacy and system vulnerability concerns. As the IoT proliferates, lack of adequate cybersecurity controls is creating billions of vulnerable (and readily accessible) devices that are emerging as the weakest link and entry point for cybercriminals.

We also observe that in the emerging world of connected machines, the lines between the intelligent layer (the realm of AI) and IoT blur quickly. IoT technology deployed in structured environments (for example, a manufacturing unit) create large sets of annotated data, which is not only the building block for machine learning and AI, but also forms important feedback loops to the physical world.

In our previous reports, we compared the key IIoT metrics of the time against the formation year of McRock in 2012. Looking back at the numbers and industry developments over the past few years, it is evident that IIoT is bigger and more ubiquitous than ever:

Industrial IoT Today

- **23.1 billion things connected to the Internet today versus 8.7 billion in 2012;**
- **16x growth in sensors shipped;**
- **15 billion machine-to-machine connections (M2M) in place today versus 2.6 billion connections;**
- **3x growth in M2M service revenue**
- **$2.0 billion in predictive maintenance revenue versus $400 million**

- McRock Capital

The Intelligence Layer

Industrial Analytics and Artificial Intelligence

As McRock continues to encounter more companies attempting scale in the intelligence layer, we reflect on some of the key themes in prevalence.

A. Brownfield "Incremental" and the Business Outcome Approach

Recognizing that most companies in this space work with brownfield deployment, the value addition is incremental (leading to transformational) to legacy systems. Clients seek partners delivering technology as only one part of the picture. An end-to-end solution encompasses industry expertise (sometimes through a mixture of in-house and external professional services) and the enablement of business outcomes.

> "mnubo's vision is really about enabling business outcomes and transformation from insight. In our initial years, we did a lot of technology sale and evangelisation. However, if you can't prove an ROI, a cost-saving, or a revenue generation then it's hard to stick."
>
> \- Fred Bastien, mnubo

The majority of start-ups in this space currently follow a consultative process to understand client requirements and determine the best approach to be taken to solve pain points. Service teams frequently follow up predictive maintenance reports for asset monitoring. We are still some time away from the full integration of AI/ML capabilities for proactive servicing in a business environment.

The Intelligence Layer: Integrating the Elements

Analytics & Feedback	*Dashboards + algorithms are customized for client requirements*
Interoperability	*Interoperability of data/sensors with legacy systems*
Middleware	*Different streams of data are normalized*
Sensor Management	*Pre-process volumes of data and filter out the noise*

A professional services driven customization layer is still an integral part of ensuring business outcomes

B. The infrastructure: Building a robust data environment

Most current analytics solutions are data-hungry. Exponential growth in the amount of data creates new opportunities but also brings additional challenges to the industry, since we need to handle a variety of data sources, both from legacy systems and newly embedded sensors. The idea and implementation of data lakes make a lot of sense as the "lake" captures three critical dimensions of data at low-cost, and in a scalable way. It is a common belief that the more data aggregated, the better, and eventually, all data will have potential analytical value. However, collecting a significant amount of raw data without strategic planning creates more noise than value, leading to data swamps and the need for an army of data scientists and engineers to perform relevant statistical analysis. Data quality becomes a new problem for data analytics vendors, as they struggle through an ocean of data from many different sources.

> ".. everybody wants a data lake, but they don't know why and they don't know the cost. In some cases companies are moving 20,000- 30,000 data values with 20 years of historian into a data lake with a prayer. Use cases have to be better defined and narrowed to the assets you are trying to fix."
>
> \- Keith Flynn, AspenTech

In a seemingly connected world, do we finally have sufficient data to make smart decisions? Although enterprises are leveraging data from their existing systems, they still have plenty of disconnected objects in production lines. In order to get the right analytic solutions, companies need to equip with additional connectivity tools (either built in-house or provided by a third-party vendor). The question could alternatively be whether enterprises have the right connectivity and are collecting the right data?

Overall, the industry tends to presume that AI is the silver bullet for the connectivity problem. If data is the cornerstone of analytic applications, then building a robust data environment is as important as developing advanced analytics and AI engines. As we go further into the Industrial IoT revolution, new problems in the way we collect, store and use data will arise and thus require the continuous attention of stakeholders in seeking additional solutions.

Market Map: Industrial Analytics Ecosystem

C. Human in the loop: Supervised data sets and leveraging domain expertise

We've come across "supervised" and "un-supervised" learning in a number of conversations with regard to machine learning and analytics. Given that "supervised" algorithms need to have data labels of historical working records from that machine, the approach requires significant involvement of engineers and data scientists, resulting in a long deployment process. As a consequence of the painful experience with the "supervised" analysis, the industry is now trending towards the idea of unsupervised solutions, which can operate without the need to understand an asset's design and models.

However, in some spaces where there are limited sets of time series data, industry experts would argue that this approach may not work, as these systems need to consume a significant amount of data. In the case that they do, the integrity of the analysis may be compromised if the data is not sufficiently cleansed. In other words, unsupervised analysis encounters difficulty when dealing with unstructured data.

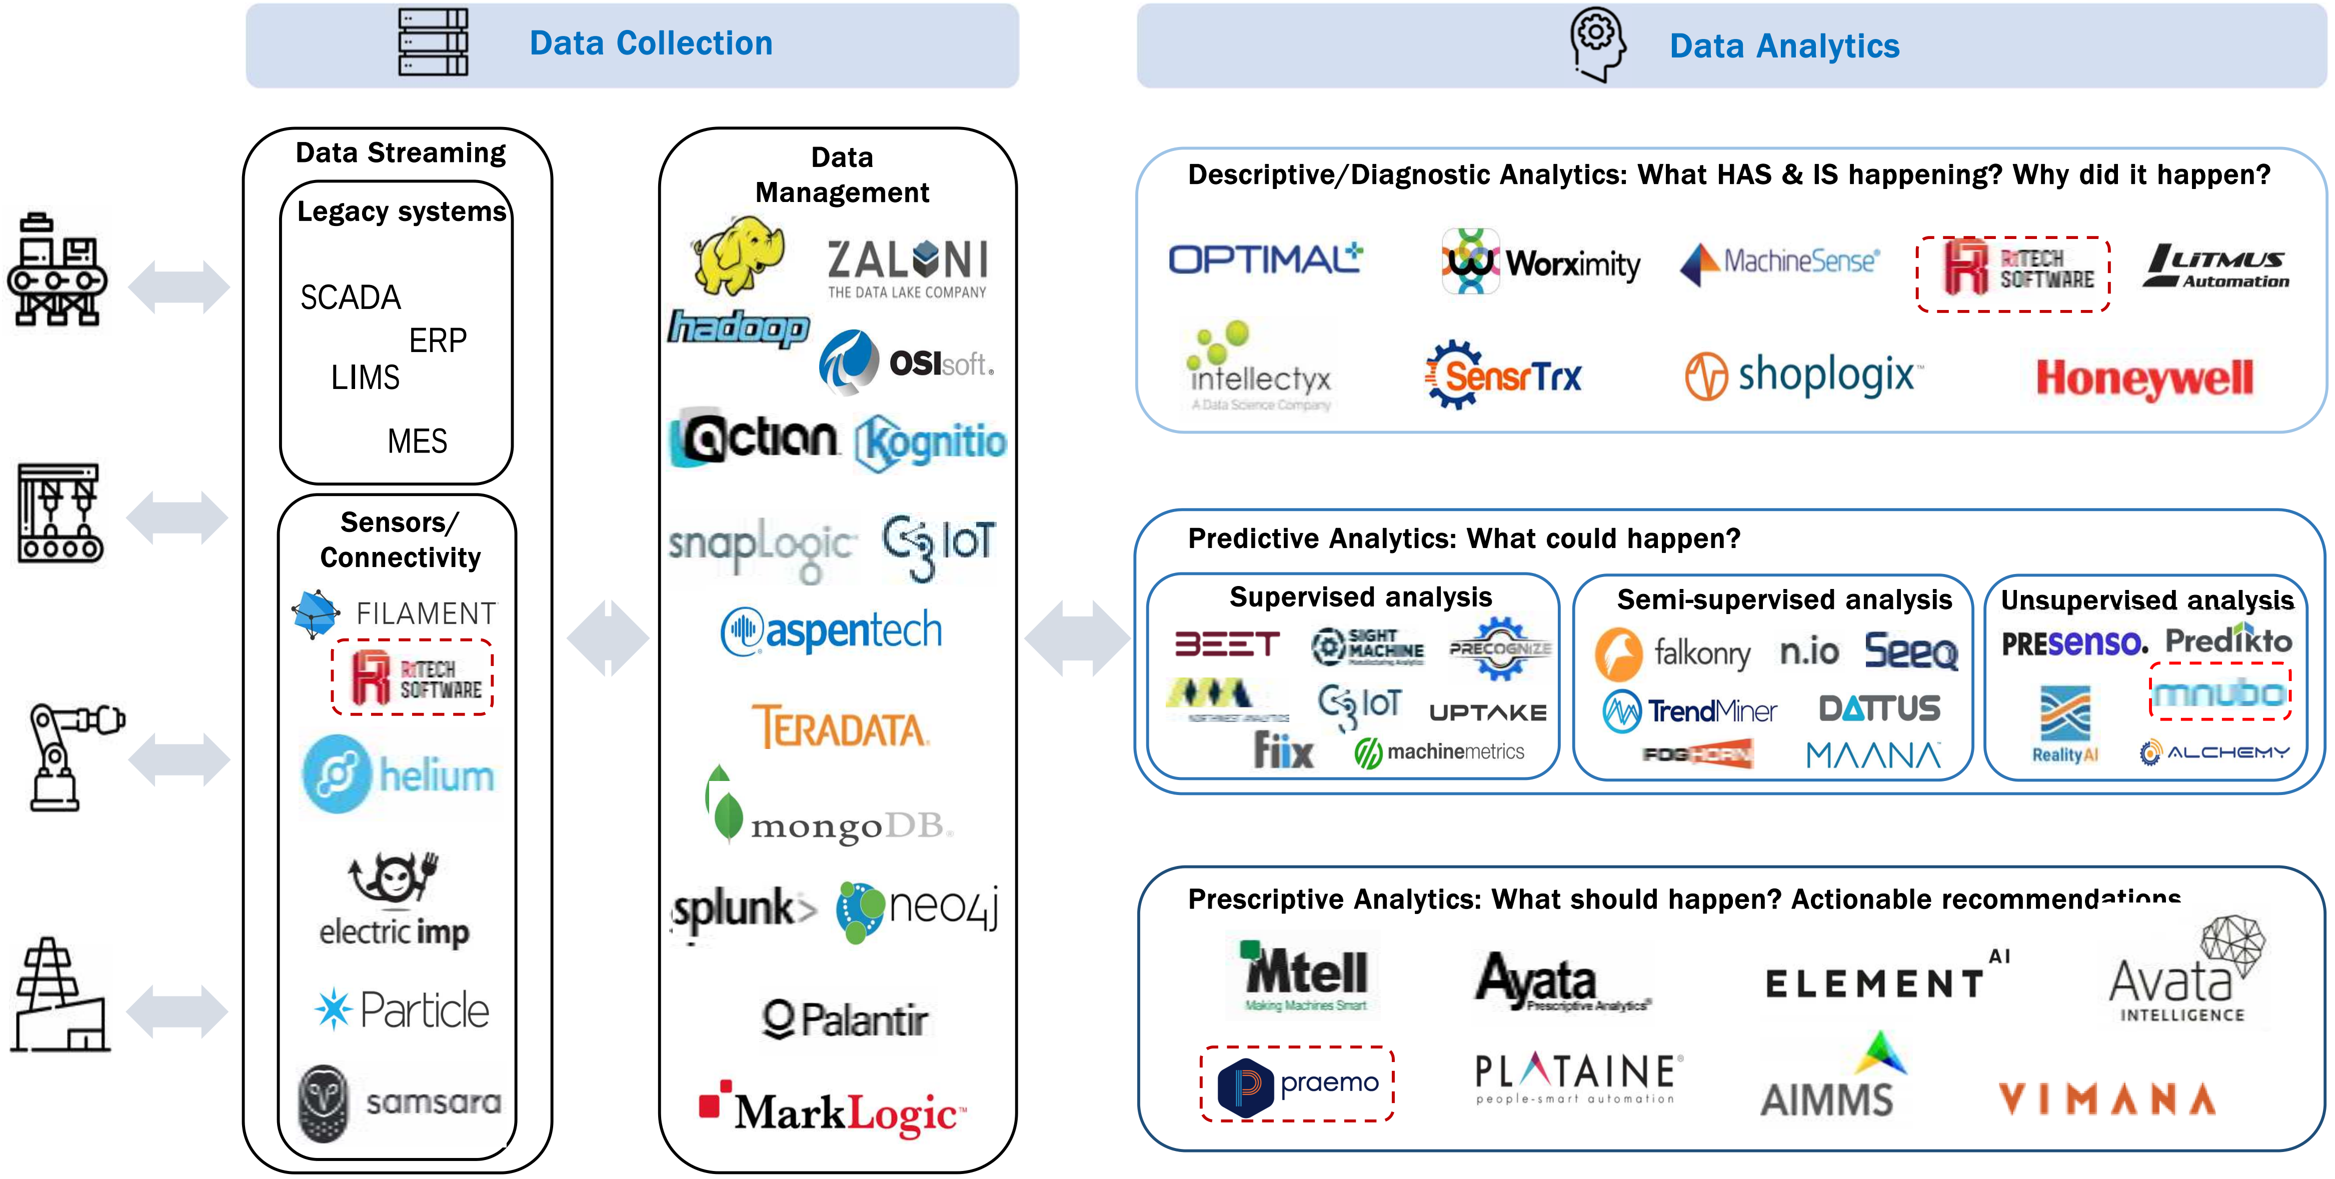

Data Collection

Data Streaming

Legacy systems

SCADA
ERP
LIMS
MES

Sensors/ Connectivity

Data Management

Data Analytics

Descriptive/Diagnostic Analytics: What HAS & IS happening? Why did it happen?

Predictive Analytics: What could happen?

Supervised analysis

Semi-supervised analysis

Unsupervised analysis

Prescriptive Analytics: What should happen? Actionable recommendations

Mtell Making Machines Smart
Ayata Prescriptive Analytics
ELEMENT AI
Avata INTELLIGENCE
praemo
PLATAINE people-smart automation
AIMMS
VIMANA

The term "descriptive" is sometimes used solely to explain what happened in the past, with real-time and diagnostic analytics categorized separately. For simplicity, we unify these categorizations

Most often, predictive analytics require the use of statistical models and machine learning techniques

Analyze the feedback coming from using the rules to take into account the actions carried out and their effects for further analysis

McRock Portfolio: Current & Exited

The journey ahead ...

> “Generally speaking, as a data scientist you want supervised data sets. However, in Industrial IoT, it’s rare that a client has a supervised data set. Consequently, solution vendors end up working with unsupervised sets all the time and hence use different techniques. The majority of use cases for semi-supervised data tend to be quite specific, sensor driven and may not always be scalable.”
>
> \- Fred Bastien, mnubo

It is important for clients and investors to look behind the buzzwords and ask the right questions about fundamental requirements for the solutions to be used in a particular industry and the scalability of the solution. Specific attention needs to be given to the technical infrastructure, resource requirement and problems to be solved.

At the current state of the technology, deep domain expertise of engineers is still the best source of reference for teaching machines and can help augment many of the highlighted deficiencies in data sets. Although McRock has come across companies which develop advanced unsupervised algorithms which can guarantee less false positives through self-learning, we still believe in the critical role of domain experts to deploy the solutions. Especially in situations where there is only a small set of data, important conclusions can be drawn from inputs of experienced engineers.

As D. Sculley of Google postulates in the whitepaper “Hidden Technical Debt in Machine Learning Systems”, ML applications are particularly prone to amplify biases in data because of subsystem entanglements and feedback loops. Given that most vendors are not reinventing the ML kernel, there is value in creating additional integrity in the enablement of data engineering, scrubbing, QA, AB testing of models. Domain driven players are uniquely positioned to offer a stronger value proposition to clients in this environment.

In the past few years, industrial analytics and AI has witnessed important breakthroughs – investors and companies alike have increasingly directed interest and investment capital to these solutions. Regardless of the hype cycle, the intelligence layer is a key component of the broader ecosystem and is critically important for creative monetization. In order to deliver long-term benefits for stakeholders, innovators need to adopt an inclusive approach based on the delivery of outcomes and the right expectations on the opportunities and the challenges ahead.

Smart Agriculture

Using data to feed the world

Agtech is a relatively small but rapidly growing vertical market of the Industrial IoT. The number of technology players targeting agriculture is increasing as investments have grown alongside some notable M&A transactions. Relative to other industrial segments, Agtech still has some path ahead before overwhelming investment interest begins.

Part of the digital transformation in the space has been driven by the next generation of farm operator/owner. Concurrently, the overall industry is also adopting smart farming as part of increasing regulation from government

Industrial IoT: Key areas of Smart Agriculture activity

- Primary data sources for insight include sensors, field scouts, field maps, commodity prices and weather. For example, field sensors allow farmers to obtain details of topography and resources in the area, as well as soil variables such as acidity and temperature. They can also access climate forecasts to predict weather patterns in the coming days and weeks
- Smartphones allow farmers to remotely monitor their equipment, crops, and livestock, as well as obtain data on their livestock feeding and produce
- Platforms are directed towards workflow automation and inventory management - utilizing data from external vendors such as agronomists, equipment vendors and grain elevators to enable richer data sets and deliver additional services
- AI and ML initiatives are focused on crop and livestock predictions as well as the minimization of data entry resulting in the efficient automation of decision making on the farm

As the industry has evolved, we believe that many companies have started with field-specific solutions and have morphed into more enterprise-wide platforms, rendering two-category approaches to the space as too simplistic.

Market Map: Agriculture technology companies and their positions in the ecosystem

As McRock continues to encounter more companies in this space, our approach to AgTech categorization within Industrial IoT is as follows:

A. Farm Management Software **- Holistic Platform and Traceability**	AgConnections FARMERS BUSINESS NETWORK AGRI-DATA SOLUTION conservis Decisive FARMING FARM AT HAND AgSquared Granular	*Farm Management Software is deployed at the enterprise level which provides farmers with a holistic view of a farm's operation. Farmers can access a variety of farm data in one place which enables them to make significant improvements in efficiency, yield and net profitability*
B. Precision Agriculture **- Improve Yield**	adapt-N MZB FarmLogs SATSHOT THE CLIMATE CORPORATION SST SMS FarmersEdge R7 Tool by WINFIELD CropMetrics MapShots semios encirca AGVERDICT Decisive FARMING VRESSON echelon Premier Crop Systems	*Precision agriculture is a field level service which improves yield through techniques such as variable rate seeding and fertilizing. Companies like Decisive Farming utilize a patented process for GIS, soil and agronomic analysis to improve yields across millions of acres of farmland* *Other horizontal analytics companies, like mnubo, are working with agriculture companies to predict crop yields using AI*
C. Crop Marketing **- Drive Revenue**	Cargill Roach Ag. Marketing, Ltd. GRAIN HEDGE FARMERS BUSINESS NETWORK FARMLEAD Decisive FARMING ADM CHS	*Crop marketing companies drive revenue for farmers by tracking near real time farm margins against the value of futures and options* *This is a highly valuable area for farmers given the margin compression and volatile commodity price environment. There are only a handful of players active in the area*
D. Telematics/Sensors **- Increase Production**	APEX SPENSA ONFARM AquaSpy cropx FARMOBILE semios FarmersEdge HORTAU Trimble AgDNA CASE IH	*This is an older segment with mature players focused on enabling increases in production by being able to sense as factors such as field moisture, nitrogen levels and pest detection* *More recently, smartphone applications and drone-based machine visioning systems have emerged to track crop health to ensure a successful growing season*

McRock Portfolio

An active transaction landscape ...

> “The opportunity is ripe for the adoption of sophisticated devices to automate data collection from primary sources and play a role in other key farm interactions. The combination of right device and right data makes for a compelling proposition … particularly with regard to autonomous vehicles application.”

- Remi Schmaltz, Decisive Farming

In many ways, farming lends itself readily to Industrial IoT given the scale of the market and it’s heavy reliance on equipment – the use of field sensors and satellite imagery for weather dates back to the early 2000’s. Low-cost sensor technologies combined with sophisticated data initiatives enable applications such as agribusiness platforms. Business Insider estimates that IoT device installations in agriculture will increase from 30 million in 2015 to 75 million in 2020, for a compound annual growth rate of 20%. There is a significant opportunity in how the data from these devices will unlock value for market participants and vendors alike.

Despite the emerging nature of the space, consolidation is also a key sectoral theme given that technology is not a particularly strong point for incumbent agriculture players and that market access is a key component of growth. A notable example is the acquisition of Granular by DuPont to accelerate its "Digital Ag Strategy". Other acquisitions include Blue River by Deere in September 2017 to advance their machine learning capabilities and the acquisition of Climate Corp by Monsanto in 2013, which subsequently purchased several other IIoT Agtech companies including HydroBio in 2017, VitalFields in 2016 and 640 Labs in late 2014.

Our view is that the macro pressures remain strong for the adoption of technology that drives efficiency and cost reductions in agriculture. The power of data and software is equally as strong in farming as it is in digital manufacturing and smart cities. The large incumbents have strong market penetration and reach and are seeking digital product extensions or up-selling opportunities. Additional market consolidation is expected given that there are only a limited number of highly credible Agtech startups.

Securing the IoT

New threats are posed by ‘botnets’ and ‘thingbots’

IoT device numbers are growing and quickly emerging as one of the most vulnerable points from an organizational security perspective. As adoption scales and newer use cases emerge, the potential for vulnerability and cyber exploitation widens.

Modus Operandi

- In the past, cyber attackers infected large volumes of private computers with malicious software and controlled these devices as a group, without the owners' knowledge. These robot networks, “botnets” in security parlance, were built in hosting environments where resources (servers, memory, address space, bandwidth) were expensive and consequently used for nefarious purposes (for example, for ransomware)
- The advent of large numbers of IoT devices have changed the resource requirements that constrained cybercriminals in the past
- For example, if attackers gain access to the administrator credentials of a manufacturer’s IoT device today, they would have instant access to potentially thousands of units, each of which are additive to their botnets - all for a comparatively small amount of work and capital
- Termed “thingbots”, IoT devices that are part of a botnet are seized fairly simply by attackers: i) Scan for vulnerable devices and apply brute-force ii) Install malware and auto-build a botnet iii) Use the botnet for attacks
- IoT devices are best known for launching Distributed Denial of Service (DDoS) attacks - attacks that are difficult for the target to defend as the traffic can be difficult to flag as malicious

> “Why wouldn’t attackers use these devices rather than costly hosting environments in which to build their botnets. With a virtually inexhaustible supply of IoT devices, we’re seeing more and more botnets, or “thingbots,” built exclusively out of IoT devices.”

- Sara Boddy, F5 Labs

F5 Labs: Timeline of Thingbot Discovery

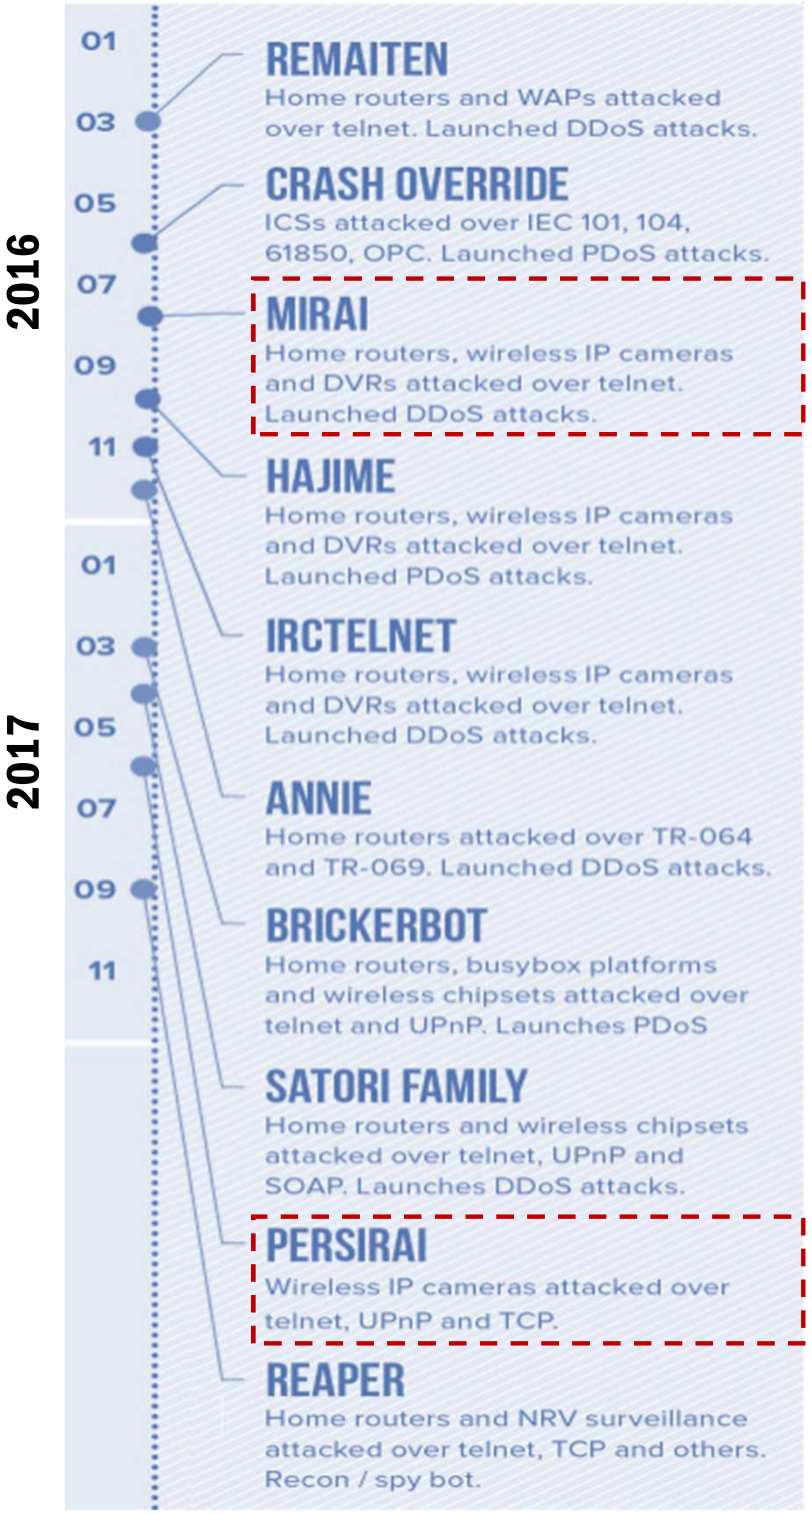

REMAITEN
Home routers and WAPs attacked over telnet. Launched DDoS attacks.

CRASH OVERRIDE
ICSs attacked over IEC 101, 104, 61850, OPC. Launched PDoS attacks.

MIRAI
Home routers, wireless IP cameras and DVRs attacked over telnet. Launched DDoS attacks.

HAJIME
Home routers, wireless IP cameras and DVRs attacked over telnet. Launched PDoS attacks.

IRCTELNET
Home routers, wireless IP cameras and DVRs attacked over telnet. Launched DDoS attacks.

ANNIE
Home routers attacked over TR-064 and TR-069. Launched DDoS attacks.

BRICKERBOT
Home routers, busybox platforms and wireless chipsets attacked over telnet and UPnP. Launches PDoS

SATORI FAMILY
Home routers and wireless chipsets attacked over telnet, UPnP and SOAP. Launches DDoS attacks.

PERSIRAI
Wireless IP cameras attacked over telnet, UPnP and TCP.

REAPER
Home routers and NRV surveillance attacked over telnet, TCP and others. Recon / spy bot.

37% of thingbot attacks are directed at non-consumer networks and infrastructure – Mirai and Persirai are the key attackers of IIoT

How Telnet creates IoT vulnerabilities: A closer look

The vulnerabilities of IoT devices emerge from the many embedded system applications in these devices. Internet routers and industrial control systems commonly leverage the remote access capabilities of a protocol called Telnet (invented in 1969) which does not encrypt communications and is an easy target for attackers. Other devices commonly connected to enterprise networks such as VoIP phones and television monitors also use the same technology.

Once attackers find an open Telnet port, they can:

- Determine what information is shared between connected devices, including the particular hardware or software model. The attacker can then exploit any known vulnerabilities associated with each
- Identify whether authentication is required. If it's not, the cybercriminal can gain unauthorized access and explore the system to see what data it contains
- Try common default accounts such as root/root, system/system, manager/manager, etc. to gain unauthorized access
- Easily perform brute-force attacks to obtain passwords for common user accounts or system (root or administrator) accounts

The use of Telnet to target IoT devices is just one more example of attackers using an older technique to compromise a new technology. IoT devices and industrial control systems present in our networks don't always get the level of security review given to a new computer server and can therefore be breached more easily

Recent spotlight enterprise attacks: Remotely networked devices were suspected to be the cause*

I. ***September 2018 – Airports:*** *Cybercriminals installed ransomware on the computer systems powering display at Bristol Airport in the UK and shut them down, forcing officials to scramble to keep the airport operational.*

II. ***August 2018 - Manufacturing:*** *TSMC, one of the largest pure-play global semiconductor companies, shut down a number of fabrication units after a 'WannaCry' malware variant spreads through the production network. The attack cost USD 170 million in lost revenue*

III. ***February 2018 – Sporting Events:*** *WiFi at the winter Olympics and display monitors were attacked*

IV. ***July 2017 – Entertainment:*** *Cybercriminals attempted to acquire data from a North American casino by using an Internet-connected fish tank. The fish tank had sensors connected to a PC that regulated the temperature, food and cleanliness of the tank*

V. ***May 2017 - Health Systems:*** *The National Health Service in the UK was infected by the 'WannaCry' malware resulting in the cancellation of 19,000 appointments and in a loss of USD 120 million*

VI. ***January 2017 & December 2015 - Power Stations:*** *Cybercriminals conducted a coordinated attack against 3 power distribution companies in Ukraine, which cut electricity for three to six hours. The attackers overwrote the firmware on the remote-terminal units that controlled substation breakers to prevent engineers from restoring power remotely. The same group is expected to be responsible for both attacks and is suspected of attempting to create a template for further global utilities attacks*

** Non Telnet based attacks*

Awareness of the IoT security risks is high on customer minds

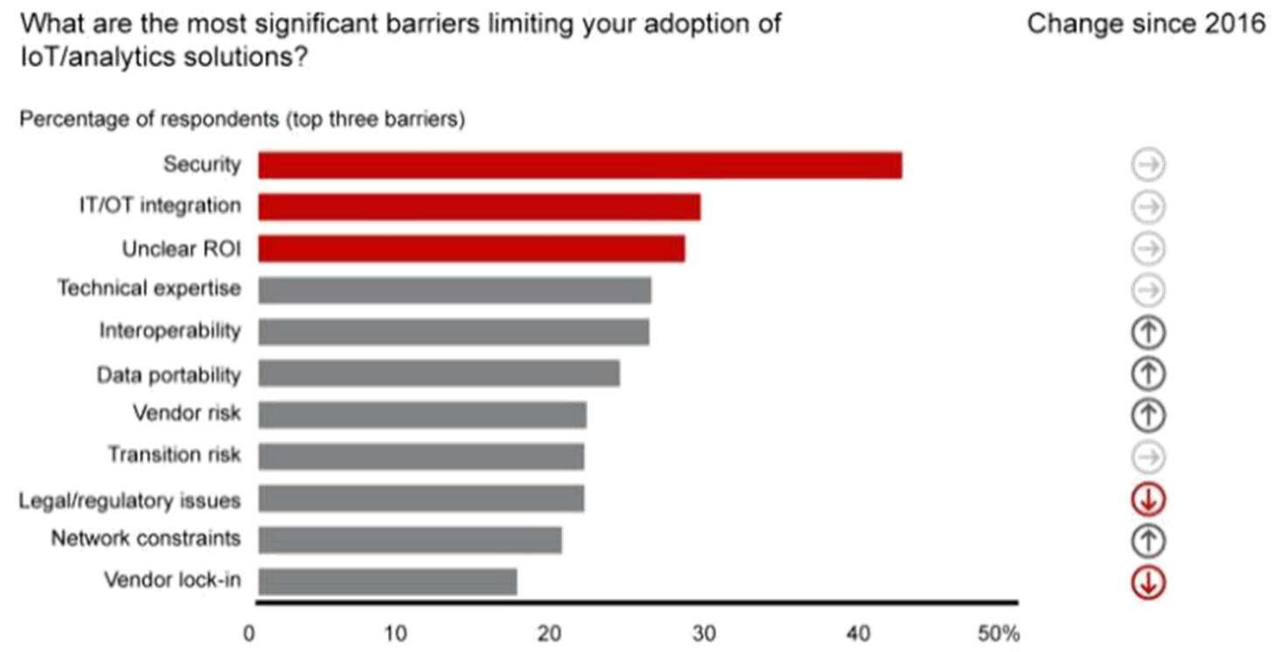

Sources: Bain IoT customer survey, 2016 (n=533); Bain IoT customer survey, 2018 (n=627); market participant interviews

The explosion of the Industrial IoT has created an imminent demand for security that was previously negligible due to the isolated nature of operational technology (OT). This demand has prompted many industries to apply a standards-based approach to the adoption of IoT technology. Due to the interconnectivity of global enterprise, a number of specific standards for IT security have been established and are enforceable. However, the OT services that control critical infrastructure largely were able to operate autonomously in the past. Additionally, OT has several priorities to concern itself with that IT does not. As a result, the standards for IT don't easily transition over to OT. Even some of the standards that have been created for the IT are simply unenforceable within the OT environment due to the diversity of each industrial environment.

As IoT devices become more complex, the vulnerabilities become larger. For example, the complex chipsets used in larger, IoT-controlled objects such as vehicles can compromise vital functions such as braking, acceleration etc. From an Industrial IoT perspective, much of the critical infrastructure used in cities, logistics, and discrete manufacturing industries is directly at risk and other enterprise networks are indirectly at risk through remotely connected devices.

Research by F5 Labs illustrates that the top 4 countries by attacks launched (either internally or against the outside world) include Spain followed by India, Russia and South Korea. China used to be a significant launch pad for attack which has declined since 2017.

While the IoT threat landscape is large, there are mitigating factors in the emergence of new security players focused on problems unique to the environment in tandem with superior DDoS strategies, greater redundancy for critical services and educating employees about the potential dangers of IoT devices.