Business Value from the Analytics-of-Things

An Executive Perspective

Richard Hackathorn, Bolder Technology

05.16 EB9406

Bolder Technology, Inc.

Table of Contents

This study examines business use cases for Internet-of-Things (IoT) applications by eight companies across several industries. The focus is: Why should business executives care about IoT? As illustrated by these companies, there is a spectrum of business motivations for implementing IoT applications, along with tradeoffs and challenges. By applying data warehousing with data analytics as the Analytics-of-Things (AoT), IoT applications can leverage business value with new business opportunities. The study concludes with recommendations to executives for properly assimilating IoT applications into their IT architecture and corporate culture.

Context

The Internet-of-Things or simply IoT is a generic term for the collection of many inexpensive sensors networked together via the Internet.[1] The central question for this study is "Why should executives care about IoT?"

For executives to care, IoT must be more than a technology that incrementally improves existing business processes. This study documents how companies utilize IoT for business use cases that reduce costs, increase efficiency, improve customer service, and exploit new business opportunities. As shown in these IoT examples, the business implications impact organizational issues and create strategic opportunities that require the attention and action of executives.

Evolution of IoT

Let's start by explaining the evolution of sensors into today's IoT. Since the 1970's, sensor technology has gradually improved and proliferated, along with shifting from analog to digital. Over these decades, three important changes have occurred.

First, sensors have recently become inexpensive and reliable as the economic barriers have faded over the last five years. A good example is the typical smart phone, which is loaded with sensors. Using people to 'sense' business events is now considered as more expensive and less reliable than sensor applications. Best business practices now suggest deploying sensors generously to measure every business asset that impacts the bottom line.

Second, sensors became interconnected. An old mechanical counter at the end of an assembly line required a person to record production data at the end of business day. Now, digital counters will report the results for each assembly line, directly into a database. Today, a single real-time dashboard monitors production flows, plus analyzing variations over the past month.

With the pervasiveness of the global Internet, the connectivity of sensors has exploded, thus the label of "Internet-of-Things". This term can be misleading since most companies utilize secure intranets for sensor data flows, instead of the public Internet.

Third, sensors are now smart. Local data storage can capture measurements and report hourly. And a single device can contain a multi-sensor array measuring temperature, humidity, pressure, video, sound and even GPS location. Smart sensors perform both monitoring and control of the device. A good example is the smart thermostat Nest[2] that senses room activity, and adjusts room heating/cooling.

Around this technology, unique business cultures with its language, expertise, and legends have also formed. For instance, the oil and gas industry has used seismic sensors to discover and exploit oil reservoirs for many decades, along which arose unique practices and expertise.

IoT Value Depends on AoT

As illustrated by the companies' experiences below, business value from IoT technology depends on the Analytics-of-Things[3] (AoT), which is the integration, management and analytics of the IoT sensor data required to distill business insights. In the past, IoT generated a data stream of measurements that drove descriptive dashboards for operational control. Today, IoT business value also comes from a holistic view of the business environment across time, locations, and organizations, along with business insights from predictive analytics.

The good news is many organizations have a data warehouse and analytic skills, implying that AoT is just focusing existing skills and tools at new sensor datasets. Further, recent advances have enabled the capture and curation of most sensor data into a low-cost data lake and linked into an integrated data warehouse, enabling innovative cross-organizational applications. Data analytics can now be applied to this integrated sensor data to extend beyond operational efficiency with new strategic services.

The next section explores a series of vignettes about how eight companies are using IoT to support specific business use cases and are evolving toward AoT. In the final sections, these vignettes are summarized into recurring themes and practical recommendations for executives about applying IoT/AoT in their company.

Company		Use Cases
A	Vehicle Manufacturer	Unnecessary maintenance, optional replacement, fleet management
B	Smart City	Electrical consumption, transformer maintenance, grid balancing
C	Electronics Manufacturer	Assembly line defective detection, reducing product returns
D	Gas/Electric Utility	Peak period pricing, electric-vs-gas heating
E	Oil/Gas Producer	Higher oil reservoir recovery
F	Railroad Company	Steel wheel failure, scheduled maintenance
G	Paper Machine Manufacturer	Process control, improved replacement parts, fleet management
H	Automobile Manufacturer	Connected car, slippery road alert, zero fatal accidents

Experiences

This section documents the IoT experiences of seven large corporations, plus a city government, based on public articles and anonymous interviews with knowledgeable persons, as listed in the table below. They are a potpourri of companies of various sizes and industries. The common characteristic is that all are Teradata customers.

Browse each of these vignettes, and imagine how you could use IoT in similar ways at your company. Note how each company has matured its IoT use cases, as it evolves toward AoT.

A – Vehicle Manufacturer

Company A manufactures and operates large complex vehicles with critical life and safety issues. The equipment is highly computerized with sensors, monitoring every aspect of its operation. Even in manufacturing, these high-value parts and tools have RFID sensors attached to track their location.

Vehicle Maintenance

The initial use case for IoT sensors at the company is to manage the daily maintenance of individual vehicles purchased by customers. By monitoring whether the sensor readings were within specified limits, the company identifies specific components for repair or replacement.

Sensors on each vehicle generate several thousand measurements, amounting to a total of a few megabytes each day of operation. The sensor data streams back to the company in small batches throughout the day. All of the sensor data is stored inexpensively on a Hadoop cluster. Via Teradata QueryGrid, the entire sensor history is selectively accessed for analysis from the corporate data warehouse.

Analytics are then applied and combined with reference data from other business functions. The analysis uses the temporal SQL extensions on the Teradata platform, which dissects the complex time sequencing of sensor data. By analyzing the interactions among sensor readings to millisecond accuracy, new analysis opportunities emerged for the company.

Unnecessary Maintenance

An extension of this use case is the reliable detection of unnecessary maintenance with predictive analytics. The company found a high percentage of false positives where vehicles were serviced unnecessarily according to the actual maintenance experience. Sensor readings indicates defects, but the parts were still in good condition. The company has improved their sensor analytics, reducing unnecessary maintenance costs while maintaining safety. When a vehicle is pulled from service, sensor data also suggest optional replacements that are likely to reduce future downtime.

> The company has improved their sensor analytics, reducing unnecessary maintenance costs while maintaining safety.

Fleet Management

Another use case is fleet management, which uses operational sensor data and equipment maintenance history to extend analytics so that managing and maintaining equipment can be optimized across all fleet customers. This has enabled the company to extend its offerings beyond just selling to renting the vehicle to their customers. Enabled by predictive analytics on sensor data, this is a new "power-by-the-hour" business model. Customers now have the option of purchasing a specific number of usage hours with a specific probability of uptime.

Efficient Custom Manufacturing

The final use case is the use of RFID sensors to automatically track valuable assets on the manufacturing floor. Most vehicles are customized and hand-crafted, which requires tracking of expensive tools and high-value parts, along with forecasting process variations that would cause supply chain disruptions. For instance, if hundreds of parts are ordered but delayed in assembly, where are those parts located when the assembly line is ready for their assembly?

RFID data was initially managed as a sneaker network, where individual files for specific part locations were shared on a need-to-know basis. Enabled by extra storage and connectivity between different data platforms, the company brought the RFID data together on their Hadoop system, allowing for easy data sharing and analytical processing.

B – Smart City

A major regional government sponsored the development of a new suburban city for twenty thousand persons along with an equivalent number of workplaces. They designed and constructed this *smart city* where all services are efficiently provided and maintained. Smart city projects usually retrofit existing buildings; this project is constructing everything new on open land, applying the latest innovations.

The initial focus is energy management with sensors monitoring building energy consumption, the low-voltage local grid, and how weather affects the entire plan. Accurate weather predictions can help to optimize heating and cooling via smart thermostats and pricing plans. It also influences maintenance and operation schedules and helps plan for disruptions, rather than just reacting to severe events.

Key objectives at this stage are energy efficiency and CO_2 emission reduction throughout the city. Additionally, there is careful scrutiny of how much should be invested in sensors and analytics to manage the optimal *energy storage facility* (batteries) for heating and cooling.

Over time, the city will expand these 'smarts' across many more services, such as healthcare, integrated mass transportation, and other opportunities to better serve their citizens.

Data-Driven Energy Management

Current infrastructure consists of smart buildings, able to optimize local energy consumption and a smart electricity grid. Both are supported by analysis and control from smart information systems with a data warehouse at the heart of operations. For example, a major asset in any electrical distribution network is the transformer. A modern transformer bristles with electrical, temperature and vibration sensors. Using data from these sensors integrated with other smart city data, operators learn when and how to efficiently maintain transformers by establishing performance benchmarks and positive-feedback loops to improve maintenance procedures over time.

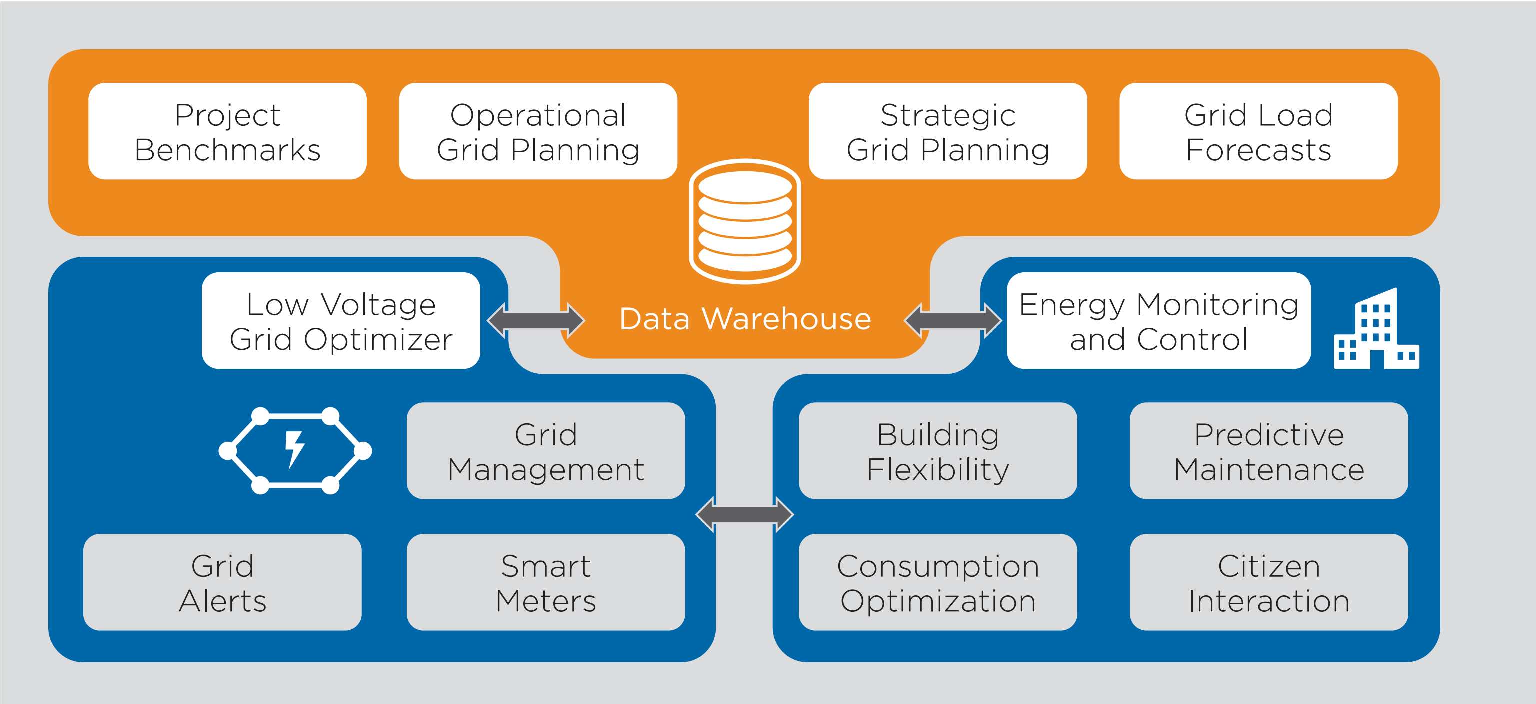

Delivering Value

Within this infrastructure, there are ten energy use cases clustered into three groups: low-voltage grid optimization, energy monitoring & control, and the Teradata Data Warehouse infrastructure. This includes benchmarking the smart city operational grid plus strategic grid load forecasting. Situation-specific analyses include pre-storm planning, actual outage events and responses, post-storm restoration, and planned-versus-actual storm disturbances and costs.

An essential component for supporting all of these use cases is a logical data model (LDM) integrating the three areas into a coherent infrastructure within the data warehouse.

For example, early results from sensors indicated that the low voltage grid (carrying 220VAC to residences and businesses) was asymmetric. Some power lines were powering most of the heat pumps, causing some lines to be heavily loaded. By rearranging the low voltage grid lines, electric usage was better balanced.

An essential component for supporting all of these use cases is a logical data model (LDM) integrating the three areas into a coherent infrastructure within the data warehouse. As an illustration of the data complexity required, the model consists of around a hundred subject areas, a thousand entities, six thousand attributes, and of course several thousand relationships.

Operational and Strategic

Sensor data is integrated second-by-second for operational management as well as year-by-year for strategic planning. City planners are establishing a set of benchmarks that enable operational performance monitoring and annual strategic goals.

Typically, such a broad scope is partitioned across independent systems with separate teams or departments having differing responsibilities, leading to severe complications balancing short-term with long-term goals. With the smart city's integrated approach, these complications are no longer an issue. For example, hour-by-hour building electrical consumption data is co-located with strategic predictive maintenance data, allowing real insight, derived from all of the data, at any time.

Early Results

In early 2016, citizens are occupying apartments in this new smart city. Data is flowing into the data warehouse, and analytics are underway. The city expects that predictive maintenance and network planning are likely to be the first use-cases to show early benefits.

Of course, even with an entirely new smart city development, there can be challenges integrating data from legacy silo systems or even with cultural differences between participating bodies. As in many traditional organizations, it is sometimes difficult to share data and adopt new practices among departments with fixed established responsibilities. One best practice they are taking is a “showcase approach” where a variety of citizens are well informed, provide feedback, and are engaged in the success of the project.

The adoption of data-driven smart city analytics is a relatively new discipline, but one with huge potential. This smart city development is already delivering on that potential in energy management. The smart city has generated real excitement for future phases, delivering new benefits to citizens and city stakeholders alike.

C – Electronics Manufacturer

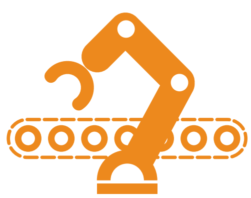

Company C is an electronics manufacturing company that uses testing devices to detect potential faults by collecting second-by-second sensor readings along the assembly line. When product quality declines past certain limits, the assembly line is stopped and corrected.

Shipping Defective Product

For over a decade, this technology has been in place to prevent defective products from shipping. However, several problems have emerged over time.

The company was using all the sensor data operationally but only archiving selected subsets. Historical analysis was not possible at a time when their products were becoming more complex. They felt that, if they could apply analytics to all the data, their failure prevention could improve. They decided to use Hadoop to capture the entire sensor stream of over ten gigabytes per hour per assembly line. However, it was difficult for business users to query the Hadoop data within their response time goals.

By using the Teradata Data Warehouse with Hadoop, the company was able to access data from their data lake. The Teradata Database analyzed the sensor data, generating reports and dashboards for a wide spectrum of business users. The response time for this information dropped from daily to minutes.

Reducing Defective Returns

Despite these efforts, defective products were still being returned by its customers, resulting in warranty costs. It was still difficult to pinpoint the manufacturing root causes for product flaws.

In the past, there were just a few manufacturing locations, making it easy for the company to manage its data integration. However, manufacturing expanded globally across more than twenty locations. Thus, procedures for data integration became more error prone and resource intensive.

The twenty assembly lines were generating about 30,000 files per day with an average size of 20 MB. It was impractical to retain these datasets indefinitely and extremely difficult to analyze the partial data that was available. By retaining all manufacturing data in a Hadoop data lake, the company is now able to quickly investigate the complete manufacturing history when a product was returned. They were now able to narrow the scope of the recalls and focus on the end-to-end manufacturing situation that produced those defects.

Think Big, a Teradata company, assisted the company with their Hadoop platform that grew from a few nodes to more than 100 nodes. In addition, they created an ingestion pipeline that brought data from dozens of systems (asset management, shipments, returns, third party contractor information), along with various manufacturing data into this data lake. Think Big implemented a combination of dataset pipeline schedulers with distributed buffer servers at each data source location. The buffer servers package and reliably transport the 30,000 source datasets to the data lake daily.

Benefits

The company reported that their product quality has improved, production yield has increased, and manufacturing costs were reduced, along with improved customer satisfaction. In addition, the return on investment for the data lake was recovered rapidly. With this solution in place, the company can rapidly trace the root cause for specific defects back to specific manufacturing changes, which are incrementally refined.

D – Gas and Electric Utility

The company provides gas and electrical service as a regulated utility to millions of residences and businesses within their diverse region. They have several customer programs as part of demand side systems.

A utility that both generates and distributes energy has the incentive to encourage its customers to use more energy. However, state regulations have decoupled their revenues from generation and provided incentives for demand reductions and energy efficiency by both residences and businesses. Hence, the utility has been actively influencing its customers to reduce their energy demand. Over the past three decades, the utility has successfully flattened the growth rate in per capita electricity consumption.

Over the past three decades, the utility has successfully flattened the growth rate in per capita electricity consumption.

The utility's energy efficiency program depends on analytics to guide various aspects of energy efficiency, such as better equipment investments, accurate load forecasting for bidding, customer insights to reduce energy usage, and reduced penalties from timely compliance reporting. The data for the analytics is supplied by many systems, including a new generation of sensors in their smart grid and smart meters.

One use case involves an alternative pricing program that bills at a higher rate during peak periods. If the utility could reduce peak demand, less investment is required to purchase additional generation capacity.

Through analytics, they determined that 40% of residential customers consume 80% of the residential electricity. Also, it was determined that the graph of the energy usage throughout the day varies considerably among customers. Analytics was able to segment customers according to their usage, such as early morning peak use versus afternoon, and so on. For each customer segment, the utility is incenting customers to change usage patterns that reduce peak energy demand and thus lower their energy costs. The same daily analysis is generalized to weeks and months, to spot annual usage patterns.

Another use case involved heating as the major energy consumer for most customers. The company analyzed the demand from electric-only customers, compared to electric-and-gas customers. The former consumed two to three times more electricity on cold days. Shifting certain customers to gas heating reduces peak demand while reducing energy costs for the customers.

E – Oil and Gas Producer

Company E is a major upstream oil and gas producer operating offshore drilling platforms across several reservoirs. A platform is like a self-contained factory out in the ocean. Platform operations have sensors "everywhere monitoring and controlling everything". For example, sensors track temperature and pressure in the wellbore and gauges controlling the flow of liquids. Sensors monitor the flow of oil from deep underground all the way through to distribution. Many platforms are unmanned. Hence operations are controlled remotely, which is safer and cheaper.

The success of the business depends on investments in two areas: First is 4D seismic analysis that tracks oil deposits over time as they are depleted. Second is permanent reservoir monitoring (PRM) that uses a seismic sensor array installed on the seabed over the reservoir.

Seismic surveys usually use specialized ships that tow many kilometers of sensor arrays over the reservoir, while creating seismic waves to map the geological structures. These survey ships must deal with turbulent ocean conditions that make precise GPS positioning difficult.

Platform operations have sensors everywhere monitoring and controlling everything.

Once an oil reservoir is developed, seismic surveys monitor changes in the oil deposits every three to six months. Only 35% of the oil in a typical reservoir is extracted because oil flows among rock pockets in complex ways. The oil flows, like weather systems, are hard to simulate even with massive compute power. Periodic surveys enable the company to achieve higher recovery rates, up to 60%. It is a tricky business, occurring miles below the sea.

Business Value

PRM is the newer practice of using a permanent seismic sensor array over the reservoir. With PRM, a survey ship only needs to create seismic waves at precise locations on the ocean surface. The business value comes from the reduced cost of repeated seismic surveys, improved accuracy of the survey measurements, and additional revenue from better recovery rates. The cost of a PRM installation can top $100M, but if a reservoir requires more than ten surveys in its life, this PRM cost is justified. Improved survey accuracy also reduces the probability of drilling in the wrong place, which is a huge financial loss. However, missing a large oil deposit is a larger loss which also justifies PRM investments.

Reservoir Data Warehouse

The challenge is getting the data out of the offshore environment for analysis. Seismic data is highly valuable and proprietary. Data for each functional group is curated and managed separately, limiting data integration. To consolidate the data across wells and surveys, data scientists from Teradata Consulting Services developed a Reservoir Data Warehouse (RDW), so that analytics can now be performed across many versions of sensor data. The RDW data integration enables a holistic understanding of the reservoir that has been difficult to achieve in the past.

F – Railroad Company

For over 150 years, this company has operated a large railroad network providing a critical transportation backbone for the country. With 8,500 locomotives and forty thousand employees, the company has continued its history of innovation with AoT applications that monitor, analyze and optimize major aspects of its operation.

Steel Wheels Failure

The first use case involves the daily collection of twenty million sensor readings on boxcars and matches each reading to specific wheels to spot any problems. For over a decade, this single AoT application has enabled the company to predict the failure of a wheel on a mile-long train traveling at seventy miles per hour, thereby avoiding the safety and cost issues of derailment.

Since overheating wheels are a failure indicator, infrared sensors are placed every 20 miles on its network of thirty thousand miles of track to monitor the temperature of each wheel. In addition, there are trackside microphones to detect grinding sounds in the wheels. For the heavier trains, the company takes ultrasound images to detect flaws inside wheels. This sensor network is linked back to the company's data center via fiber-optic cables laid along the tracks.

Within five minutes of the sensor reading, pattern-matching analytics flag outliers. This enables operators to pull a train from service or to slow the train for repair at the next station.

> The sensor data analysis has reduced wheel bearing derailments by 75%.

The company is planning to replace their entire IoT infrastructure with an improved bearing-failure system. Sensors will be embedded directly into the wheels, bearings, and other critical components to collect data wirelessly. This will enable better prediction of wheel failure weeks in advance, avoiding the congestion of trains stopped mid-route.

Strategic Growth Goal

The strategic goal is to grow its business by optimizing freight volumes through a limited number of locomotives and railcars. The company realizes that this goal requires a constant balance between analytics-driven automation and human judgment. The key metric for this growth goal is velocity—the amount of freight moving through the track network in a giant pipeline. All information funnels into a main dispatch center that displays real-time information on tracks, trains, weather, fuel consumption, and repairs at its one hundred rail yards.

The company realizes that this goal requires a constant balance between analytics-driven automation and human judgment.

The dispatch center is where the minute-by-minute operations blend with the weekly tactical plans striving toward strategic goals. The twenty-million wheel readings usually indicate fifty hot-wheel problems daily for which a train might be stopped immediately if serious enough. Those trains are checked at the next terminal or at final destination. Stoppages ripple throughout the network, much like airline cancellations.

Bottom-Line

Over the past thirty years, sensors and analytics helped U.S. railroads reduce freight costs by half in inflation-adjusted dollars. In addition, fuel miles per ton-gallon has also halved over that period. That benefits everyone.

G – Paper Machine Manufacturer

Company G manufactures large machines for producing tissue and paper products, along with automation systems and services to support their customers' operations throughout the world. Their customers are large manufacturers of consumer goods. To increase revenues, the company competes to sell long-term service contracts that insure the machines work at optimum efficiency.

Business Problem

It is critical for the paper machines to operate efficiently with low maintenance cost and long durations between shut downs. Any downtime is a major concern, whether scheduled or not. Scheduled maintenance usually occurs every few weeks. Consumable parts (like felts, rollers, bearings) are replaced during every shutdown. Accurate prediction of parts' wear-out allows the operator to delay scheduled shutdowns.

Improving Scheduled Maintenance

Each paper machine has over ten thousand sensors to measure speeds, pressure and vibration on reels and rollers. Some sensors detect an impending failure and initiate a quick unscheduled shutdown before any damage occurs. Other sensor data is analyzed to improve scheduled maintenance. The data is collected in batches from each machine daily (real-time data is not required). Today, Aster nPath predicts how long a certain machine can operate on its consumables before a shutdown is required.

Improving Consumables

The sensor data is also used by the company's R&D group to build better consumables. Cheaper parts from China are displacing some of the company's consumable sales. However, via analytics, the company's higher quality parts now operate for 20% longer and are more economical. Thus, the company analyzes each customer machine individually: When does vibration in rollers wear out consumables? How fast should the machine run? How does various kinds of pulp and chemicals affect wear-out?

"By better understanding roll performance and roll maintenance needs, it is possible to extend roll service intervals by 20% and thus extend roll run times. There is no longer a need to change a roll just to make sure that it will not break."

Different roll cover materials affect roller wear. Rolls are resurfaced on a grinding machine during maintenance cycles. Combining grinding machine and paper machine data with analytics can determine the roll condition and deviations from normal. Hence, the company can extend its machine lifetime with a custom maintenance schedule and matched cover materials specifically for that paper machine.

Fleet Benchmarking

The company is integrating sensor data from several hundred factories, each operating several paper machines, each having over ten thousand sensors sending data at one second intervals. Providing benchmarking reports is a value-added service to their customers. The benchmarks show the wear pattern of specific consumables in specific machine models worldwide by usage type. Thus machine operators can choose the lowest cost, best performing parts.

The company feels that value-added service offerings, like fleet benchmarking, are critical to remaining profitable and competitive for their cost-cutting customers. These offerings are increasingly performance-based contracts to provide a low production cost per ton of produced paper products. A multi-year pilot project with Teradata Aster is proving the business value of these contracts.

Best Practices

The raw sensor files are difficult to process, especially with Python programming. The company switched to Teradata Aster Analytics to parse raw sensor data because it was much easier. Sensor data (semi-structured JSON) and maintenance data from SAP ERP software (structured) is much easier to integrate using Aster. The data scientists quickly learned how to apply the analytic methods using both SQL-MR and R. Currently a handful of data analysts generate reports with graphs as PDF files to send directly to operators of customer's paper machines.

Benefits

The company stated that their data warehouse for analytics can analyze huge data volumes and thereby enable their customers to make better maintenance decisions.

> "Technological advances enable us to utilize embedded intelligence better … allowing our customers to improve their performance by utilizing integrated data for better benchmarking, predictive models and best practices."

The company was excited about the ease of deployment for its first customer. The company is striving toward unique visibility in global markets in its Industry 4.0[5] initiatives.

H – Automobile Manufacturer

Company H is a global automobile manufacturer who has successfully created an integrated data warehouse supporting most of its business functions. To leverage their IDW investment, the company formed a cross-organization business analytics group working closely with IT. The company has established that data, analytics and innovation will be differentiators for their automotive product. It is a fundamental digital transformation, initiating innovations like autonomous driving, fleet car-sharing, auto cloud-based services, and personal service technicians.

> Their approach is combining predictive analytics, cloud computing and autonomous driving, along with extensive use of sensors.

Car Sensors Are Not New

For a decade, most auto makers have equipped their vehicles with sensors to capture diagnostic trouble code (DTC) data in the engine control units (ECU) for better vehicle maintenance. Since the data was not centrally collected and analyzed, the only person using this data has been the service technician. Now, the company is using DTC data for monitoring conditions of batteries, such as when the battery has a no-start error. This has improved the handling of battery warranty situations. Analysis also extends to customer services.

With customer permission, 80-90% of all cars are connected to headquarters. Customers are notified in advance when certain parts are wearing out and, in contrast, also when maintenance visits can be skipped. This saves both the company and the consumer time and money.

The Connected Car

The safety initiative is driving the company to combining sensor and cloud technologies so that the vehicles are always interconnected. A two-year pilot program is underway for a thousand cars that share data with each other, the company and local municipalities, at the discretion of the owner. This allows the exchange of real-time data on road and traffic conditions. For example, if the sensors in the wheels of one cars detects slippery patches, it alerts its braking system and those of nearby cars to slow down. Other sensor-based safety systems are lane keeping aids, collision warnings, and adaptive cruise control.

The company is streaming this data into the Teradata Data Warehouse to be combined with data from internal engineering, servicing dealerships, and customer interactions via on-board systems. Analysis includes how changes in rear cameras to wheel sensors are improving safety. There is better support for cross-organization projects to understand product defects and trace them to origins in manufacturing.

Zero Fatal Accidents
Their latest AoT initiative focuses on safety with the objective of eliminating fatal accidents within five years[6]. Several car models have already achieved a driver death rate of zero over recent years[7].

However, total annual fatalities in the USA across all manufacturers have remained above thirty thousand for a decade. One approach is to curb accidents by reducing human error, which is a factor in most accidents. Hence, the company is "gradually taking the human driver out of the system" to improve safety. Their approach is combining predictive analytics, cloud computing and autonomous driving, along with extensive use of sensors.

Recurring Themes

The above vignettes describe the experiences of companies adopting various AoT use cases. What were the recurring themes that motivated these companies?

Value Expansion

IoT can be like breeding rabbits—fun, easy and profitable! Once the infrastructure is established, the initial use case stimulates ideas and questions about product improvements and business opportunities. That leads to further analyses, more instrumentation, deeper questions, and new use cases.

> Company A realized… it is the beginning of an evolving infrastructure with impacts spanning many business functions, such as failure predictions, accident investigations, long-term service and maintenance strategies, design testing, and manufacturing efficiency.

For instance, Companies A, F and G all started with a single-purpose use case (like product failure detection) and evolved into scheduling maintenance, improving replacement parts, and fleet management. This last use case resulted in an innovative business value-added service offering (and revenue stream), which was possible only with a deeper understanding of the entire product lifecycle.

Hence, the initial single-purpose IoT use case is a value-expansion enabler for:

- Refining the effectiveness of the initial use case (more accuracy, finer resolution)
- Expanding the use case to other operational issues (more applicability)
- Collaboration among functional units (more linkages around related processes)
- Introducing new data-centric services (new revenue opportunities)

Why should executives care? … because IoT value expansion projects require executive leadership and support to guide efforts toward business value.

Operational to Strategic

Contrary to popular opinion, IoT is not applicable only at the operational layer. Long-term value impacts of IoT come from maturing operational IoT use cases into strategic AoT business initiatives. In other words, IoT use cases should evolve upward within this pyramid, which requires executive leadership based on a strategic vision.

For instance, Companies B (smart city), F and G have IoT use cases that have strategic policy implications. In particular, the smart city has co-located the operations team optimizing hour-by-hour building consumption with the strategic team predicting long-term maintenance program.

Why should executives care? … because for companies lead by insightful executives, understanding and exploiting AoT strategic potential can be an industry game-changer.

Technology Assimilation

With the assimilation of any technology, organizational culture plays a critical role in its successful adoption. Since many industries have used IoT technology for decades, unique 'sensor-centric' cultures have formed with their own language, expertise and legends. In particular, recommendations based on analytics are sometimes viewed suspiciously by established experts who have honed their skills over decades of experience.

For instance, Company E has seismic experts who are adjusting to new analytics using the more accurate PRM sensors. Company G employs senior machine operators with decades of experience operating large paper mills, whose practices are being 'enhanced' by predictive analytics. The challenge is to achieve a synergism that blends human and analytic expertise to achieve results that cannot be done separately.

Why should executives care? …because executives should lead cultural changes required to realize IoT business potential by showing that IoT applications are synergistic, not competitive, to human judgment.

Essentials for Driving Value

In summary, business value is contingent on a company maturing from IoT operational use cases to AoT strategic initiatives, the essential drivers of which are shown below.

Each of these IoT drivers were explained previously. However, note the crucial role of AoT drivers. Without AoT, the IoT use case is limited to a single-purpose application that monitors one business process. To unleash IoT's potential, an infrastructure should support the following:

- **Data Curation** – Clean, organize, and archive raw sensor data into a data lake or refined data into an integrated data warehouse.
- **Data Warehousing** – Integrate data as a continuous effort to combine subject areas into an cross-organizational data warehouse.

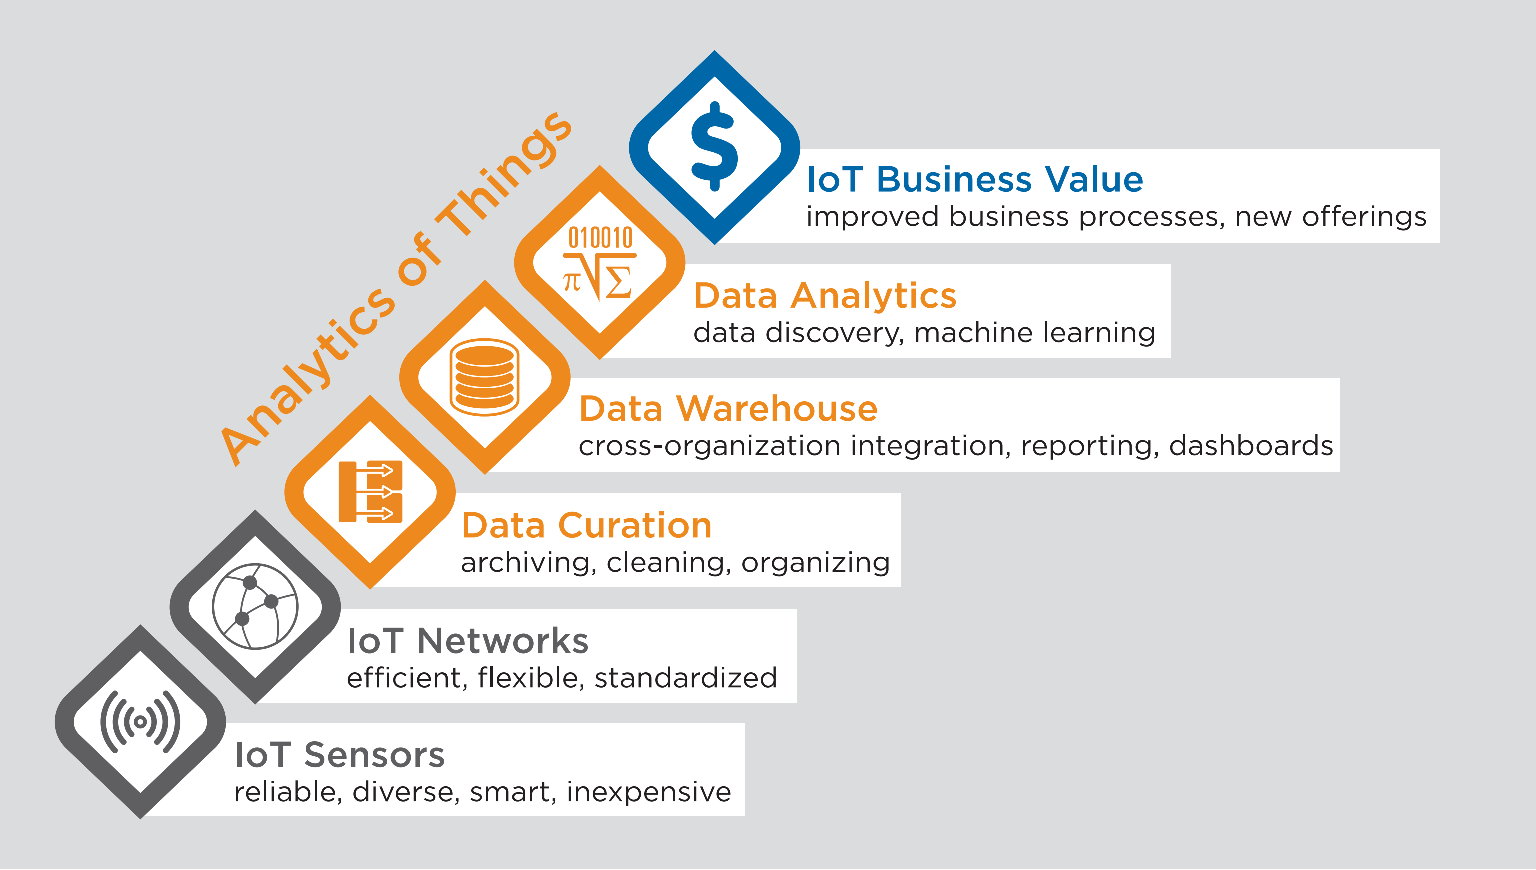

- **Data Analytics** – Establish benchmark trends, create operational/tactical dashboards, provide easy access to detailed data, govern standard test datasets, and apply predictive algorithms.

For instance, Company E is creating a seismic data warehouse combining sensor data across reservoirs and surveys. Companies B and G are using historical performance data to design and manage their new fleet management offerings. Company C is curating all their sensor data into a data lake to quickly find root cause for returned defective products.

Why should executives care? … because executives should set an evolving vision for achieving IoT business value, along with supporting essential infrastructure.

Recommendations

For executives, the following practical recommendations are suggested:

1. Be aware of IoT use cases within your company and across your industry.
2. Lead the adoption of IoT use cases based on a clear vision of business value, while avoiding technology for its own sake.
3. Challenge your company to expand IoT use cases from single-purpose, to multi-purpose, across functional areas, and toward new value-added service offerings.
4. Realize that IoT has limited business value without ongoing infrastructure investment to support AoT data integration and data analytics.

The good news is that Analytics of Things technology is rapidly improving and that existing IT infrastructure can be leveraged. The limitations to realizing its business value are lack of creativity to conceive of new approaches and the willingness to embrace those approaches.

Endnotes

1. See **https://en.wikipedia.org/wiki/Internet_of_Things** Good (but diverse) information on IoT. Note the Application section for ways that IoT is being applied across society.
2. See **https://nest.com/**
3. The term Analytics-of-Things was first suggested by Robert Stevens of First Analytics in a conversation with Tom Davenport, who popularized the term in a WSJ article at **http://blogs.wsj.com/cio/2014/01/29/the-analytics-of-things/**.
4. The verb *curate* is relevant to a data lake since its function is more than capture. A data lake should also preserve, organize, and provide context. For instance, a museum can 'capture' an ancient artifact, but its value is not realized until that artifact is properly curated by the museum.
5. See **https://en.wikipedia.org/wiki/Industry_4.0**
6. WSJ CIO Journal, Kim Nash, Volvo Works to Scale 'Death-Proof' Car Effort Using IT, February 12, 2016
7. Insurance Institute for Highway Safety, **http://www.iihs.org/iihs/topics/driver-death-rates**

About the Methodology

The methodology for this study is to listen carefully to pioneering companies in big data analytics and to report accurately on their perceptions. The intent is to contribute to professional education—to share the experiences and best practices with other IT professionals so that we can mature as an industry, amid escalating business challenges and rapidly evolving technology.

The primary author is Richard Hackathorn of Bolder Technology, who appreciates the insights shared by a variety of IT professionals. In particular, Dan Graham of Teradata deserves credit for many substantive improvements to this study.

Finally, a sincere appreciation is given to Teradata Corporation for its sponsorship in conducting this study and for permitting open and independent access to its customer community.

About Bolder Technology

Bolder Technology Inc. is a twenty-year-old consultancy focused on Business Intelligence and Data Warehousing. The founder and president is Dr. Richard Hackathorn, who has more than thirty years of experience in the Information Technology industry as a well-known industry analyst, technology innovator, and international educator. He has pioneered many innovations in database management, decision support, client-server computing, database connectivity, and data warehousing.

Richard was a member of Codd & Date Associates and Database Associates, early pioneers in relational database management systems. In 1982, he founded MicroDecisionware Inc. (MDI), one of the first vendors of database connectivity products, growing the company to 180 employees. Sybase, now part of SAP, acquired MDI in 1994. He is a member of the Boulder BI Brain Trust (BBBT). He has written three books and has taught at the Wharton School and the University of Colorado. He received his degrees from the California Institute of Technology and the University of California, Irvine.

About Our Sponsor

Teradata helps companies get more value from data than any other company. Our big data analytic solutions, integrated marketing applications, and team of experts can help your company gain a sustainable competitive advantage with data.

05.16 EB9406