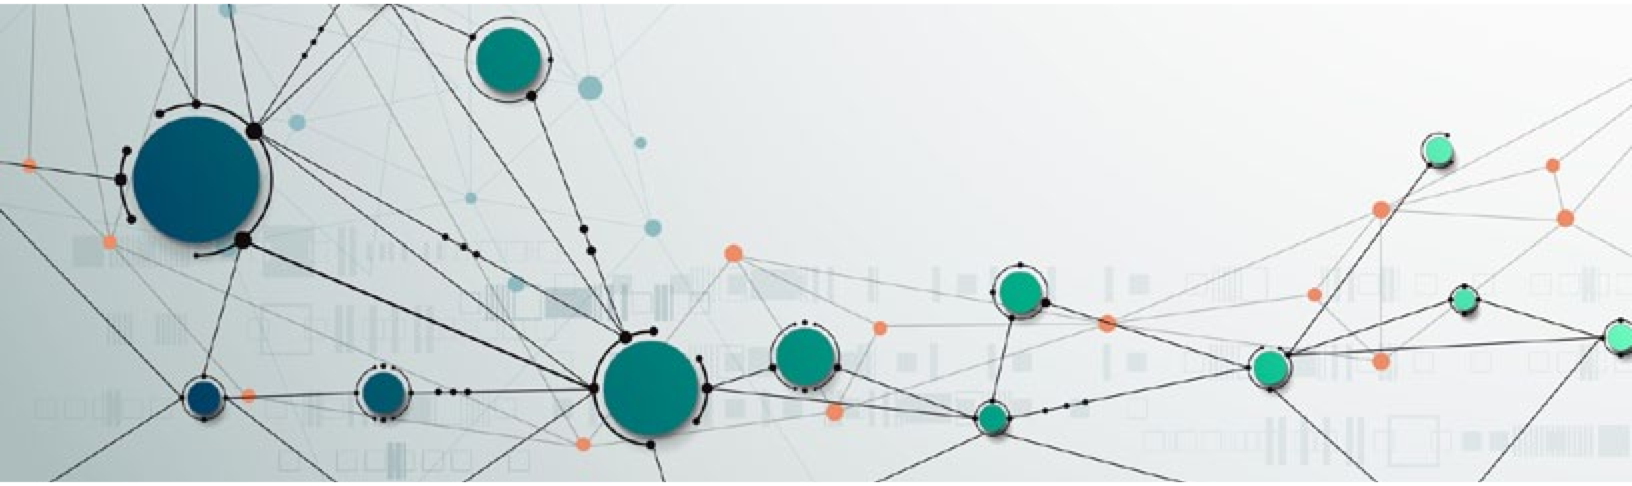

# **Edge analytics:**
# fast decisions for the IoT

eBook

RTInsights.com
Accelerate Your Business with Real-Time Insights

# Table of Contents



**Figures:**

# Why edge analytics is crucial for the IoT

Marc Benioff and the company he founded in 1999, Salesforce, called for an end to software. Customer relationship management (CRM) hosted in the cloud and subscription-based applications sprung up throughout the enterprise technology sector.

Gradually, the rest of the world fell into the cloud's orbit, with Amazon Web Services (AWS) moving the concept several steps forward by allowing anyone to host whatever they wanted in Amazon's global data center footprint. Others followed over the years and the global enterprise server market has cratered as a result. A former colleague recently told me, "If you're not talking cloud, you can't talk about technology—CIOs are now technology brokers for the business."

Every model has its strengths and its weaknesses and, like most things, a healthy mix is often the best path. With the emergence of the IoT, a mixed approach between centralized and decentralized computing models is more than just smart business: It's possibly the only way for IoT services to reach global adoption.

**BY 2018, SOME 40 PERCENT OF IOT-CREATED DATA WILL BE ACTED UPON CLOSE TO, OR AT THE EDGE OF A NETWORK.**

Estimates range that we will see 15, 30, even 70 billion connected devices globally by 2020—all generating streams of data getting pushed across the various regional networks through the internet. But not only is this unnecessary, it's often more advantageous to push analytic capability down to the outer section of the network, as close to the endpoint as possible. Rather than bring the data to a centralized hub for analysis, capitalize on opportunities at the edges.

Why is it more advantageous? Because speed enables automation—and automation is a key benefit of IoT-related projects such as continuous uptime, safety alerts and condition-based maintenance. These all rely on IoT data and will probably become highly automated as a result. If the analytics are done at the edge, then the speed at which these processes are executed is exponential. The faster you get the data, analyze the data and react, the more value is captured or value loss mitigated.

# Edge analytics: advantages and use cases

In edge analytics, devices or the gateways can handle the analysis. Advantages in doing so include:

**Reduction of network bottlenecks:** Some data, for example video employed in smart city applications such as traffic management, is so large that it could congest the network. A network with a bandwidth of 100 Gbps, for example, can support uploads of 1080p streams from only 12,000 users at YouTube's recommended upload rate of 8.5 Mbps, according to a recent article in Pervasive Computing. A million concurrent uploads would require 8.5 Tbytes per second.

**Fast response times:** Applications such as energy production from wind or solar power plants and monitoring of ill patients require response times of a minute or less. When that data is sent to a central location for analysis, it loses its value. The need for speed is so great that by 2018, some 40 percent of IoT-created data will be stored, processed, analyzed, and acted upon close to, or at the edge, of a network, according to a report by IDC Futurescape.

**APPLICATIONS THAT REQUIRE LOW LATENCY OR ARE HIGHLY DISTRIBUTED ARE PERFECT CANDIDATES FOR PUSHING ANALYTICS TO THE EDGE.**

**Data filtering:** Allows analytics to be performed on actionable data; only necessary data is analyzed or sent on for further analysis.

Applications that require low latency or are highly distributed are perfect candidates for pushing analytics to the edge. Businesses involved in pipeline monitoring, connected oil and energy grids all have high degrees of distribution and make edge analytics highly attractive. These organizations must consider the need to engineer new products, devices and sensors with the IoT in mind to fully take advantage of what edge analytics can provide.

Warehouses, logistic managers, factory sites, retail sites and hospitals are all candidates for a decentralized model for analytic computation. Event data across regional distribution locations can be analyzed to improve real-time fulfillment models. Traffic and weather data within specific regions can reduce delivery costs while improving on-time fulfillment rates.

Plus, retail locations can restock in real-time, hospitals can detect device failure, and wind farms can course-correct turbines to maximize energy capture.

Some other use cases include traffic management with smart lighting and smart parking spaces; retail sales and customer experience management with beacons, smartphones, and smart vending machines; preventive policing with smart lighting and gunshot detection technology; and congestion management and safety in smart transportation systems.

### Challenges for edge analytics

In a March 2015 survey of 203 IoT professionals conducted by Dimensional Research for ParStream, participants said the chief technological challenge for the IoT was timely collection, processing, and analysis of data.

Thirty-six percent of respondents said it was difficult to capture useful data, and 44 percent said there was simply too much data to capture effectively. "There used to be a day, many years ago, where the word 'terabyte' meant something," said Syed Hoda, chief marketing officer for ParStream. "We're well beyond that now. Terabytes of data are just a start in some of these use cases."

**VELOCITY AND VOLUME ARE CHALLENGES FOR BIG DATA ANALYTICS.**

Ninety-two percent of survey respondents, meanwhile, said they couldn't capture data fast enough. Whether data is real-time or not often depends on the use case—for wind turbines, that might be a few seconds, said Hoda.

Respondents also had trouble with software, and reported that the analysis capabilities were not flexible enough to support queries the companies wanted to ask.

## IoT architectures for edge analytics

Edge analytics requires hardware (routers, servers, gateways) as well as analytics software that can analyze data from endpoints.

It is worth noting that edge computing doesn't replace all analytics. Under a reference model for what is known as fog computing, edge devices may perform analytics that are crucial to protection and control systems, and which require machine-to-machine response times of sub-seconds. Visual analytics and processes may take a few seconds, and the cloud would update rules and patterns.

For IoT architectures, "What might happen is that you can have multiple sensors going through multiple gateways, which might send data to a cloud server where the analysis can happen," said Fern Halper, director of TDWI Research for advanced analytics. "And that might be the case where there's a lot of analytical heavy lifting to do."

"What happens is that the [analytic] model might be produced in the cloud and then instantiated in the edge device or the gateway," said Halper. Where the analysis is done, however, depends on the business case and the efficacy of the edge device or gateway.

**"ANALYTIC MODELS MIGHT BE PRODUCED IN THE CLOUD AND THEN INSTANTIATED IN THE EDGE DEVICE OR GATEWAY."**

Certain types of heavy analyses cannot be performed at the edge. And sometimes the edge device, such as a smart meter, is not capable enough to run analytics.

For instance, in the case of energy consumption, "usually smart meters on homes are not capable enough to run analytic models," said John Thompson, GM of advanced analytics at Dell. Smart meter data could be sent to a gateway, which performs data filtering, and then sends data to a server running at a city or regional level.

"We don't really need 10 million data points to say the smart lightbulb is on. We're really just interested in the state change," he said.

Thompson described an IoT architecture using Dell Statistica, but "there are other analytic products out there that have similar operational characteristics," he said. Statistica can export an analytic model anywhere on the IoT network in native languages such as Java, C, C++, SQL, and PMML.

## Figure 1: Dell device and cloud analytics

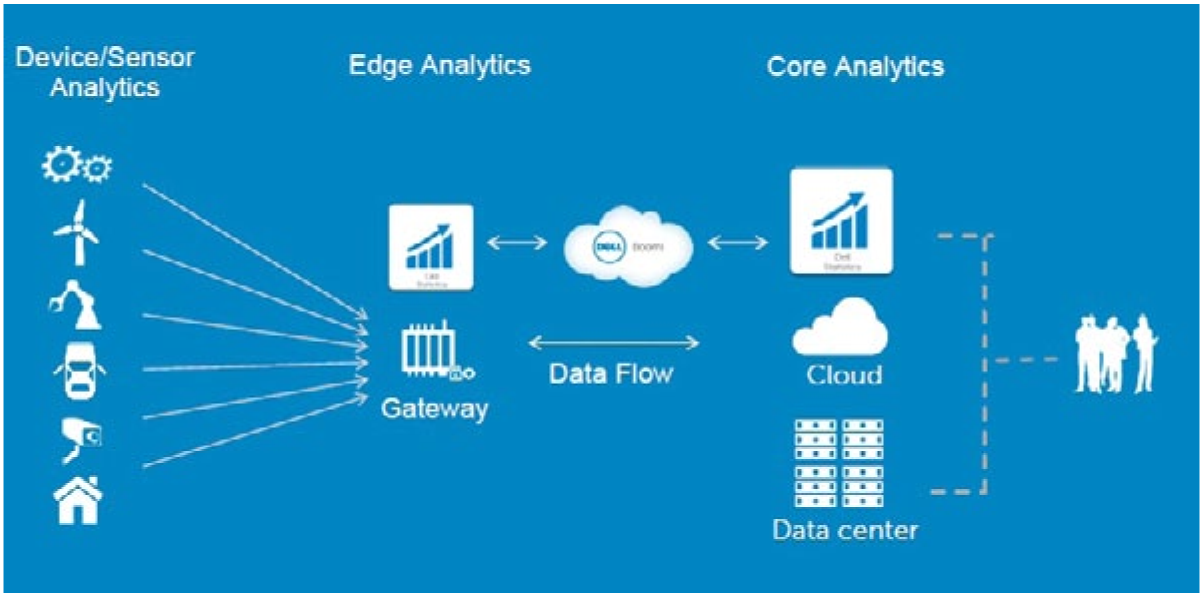

In addition, Dell Boomi, a cloud integration platform, can take the model and integrate it with data from customer-relationship management systems such as Salesforce; data from a private cloud (for instance data from overseas operations); and data from devices. The platform can also take analytic models with attached IP addresses and Boomi delivers that model to the associated IP, which might include a natural gas plant, a John Deere tractor, or a Salesforce cloud.

"We're thinking that predominantly the analytics are going to be running in those gateways," Thompson said. "Now if you had environments such as windmills, or industrial environments, or cars that are capable of running models, you can certainly push models all the way out there."

Thompson said that rather than thinking of the edge versus the center, he thinks of IoT architecture as "concentric circles in the network all the way back to the center." Each layer could have their own analytic models appropriate for the layer.

Cisco takes a similar approach with their model for "fog computing," in which edge analytics is targeted at high-speed, low-latency needs, with response times ranging from milliseconds to under a minute, while longer-term analysis takes place in the cloud:

## Figure 2: Cisco fog analytics model

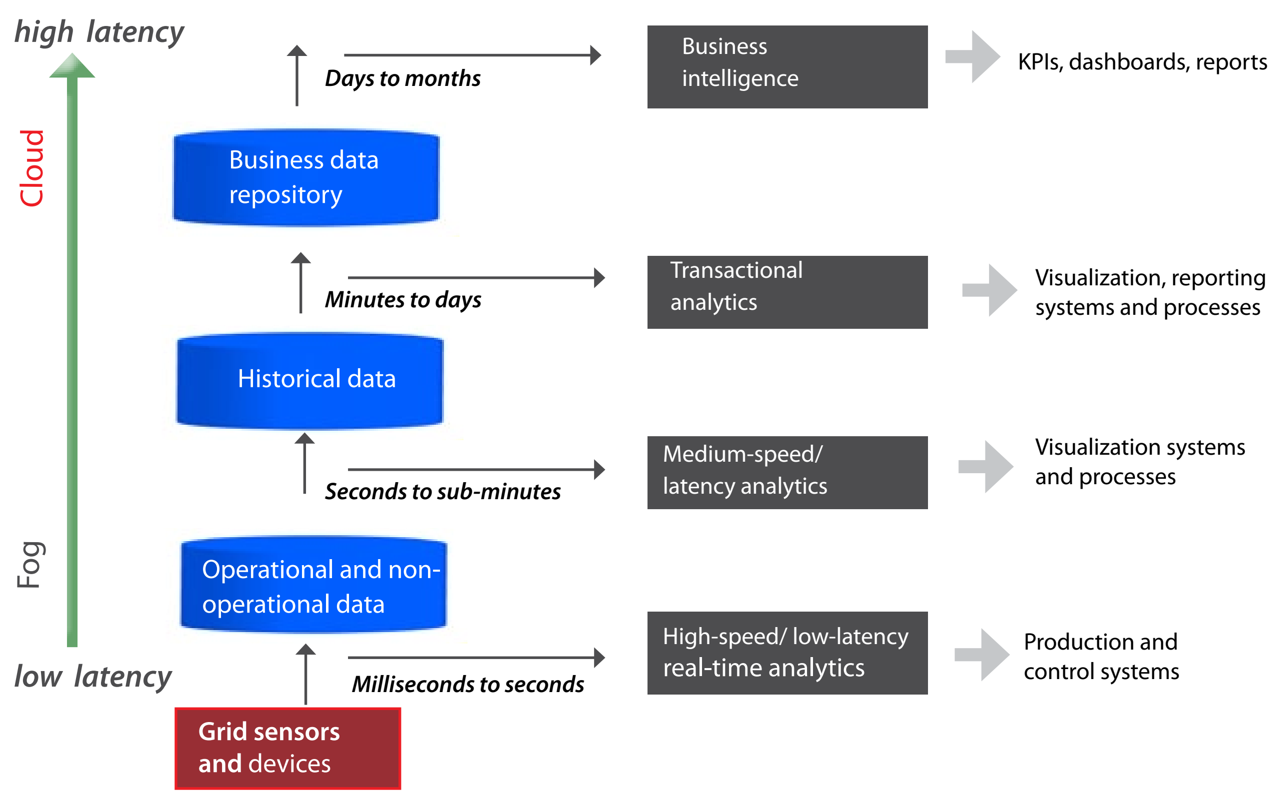

# Why gateways and controllers are critical to edge analytics

Because devices generate a tremendous amount of data, an effective IoT architecture features middle-tier gateways to manage high volumes of continuous data, giving companies only needed data.

The IoT begins with front-end sensor data and ends with back-end systems using analytics and decision-making capability to drive intelligent actions. What's missing from the IoT discussion is how to best manage the movement of this data from the front-end to the back-end. This is where the concept of a middle tier supporting controllers and gateways fits in.

**Front-End Edge Devices**. Front-end edge devices vary in intelligence and are nearly always tasked with data collection. Front-end devices therefore usually contain one or more sensors. The intelligence of front-end devices varies from simple sensors to intelligent micro-process controlled devices. Regardless of the intelligence of the front-end device, it needs to know how to identify itself, publish data or events securely, and interact with its parent gateway or controller. Smarter front-end devices will have the capability to bi-directionally interact with upstream controller and gateways and respond to directives.

**Controllers and Gateways.** Controllers and gateways are middle-tier components that primarily collect information from front-end devices. However, the role of controllers (consumer focused) and gateways (commercial focus) in consuming this information varies widely. In some cases a gateway will primarily be focused on data collection, data aggregation, and then transmission of the data to a back-end system. The periodicity associated with this data collection, aggregation, and transmission will vary from real-time streaming and queuing to batch processing. In other situations, the gateway or controller will also analyze and act on the data based on a predefined policy in order to automatically interact with front-end devices. The location of controllers and gateways will also vary based on the intelligence and constraints of the network connection to the front-end devices, environmental factors, and security concerns. Common networking between edge devices and controllers/gateways includes BT, CAN, LAN, WAN, WIFI, USB, and Internet.

**Back-End Systems**. Back-end systems will typically function as the repository for information transmitted from controllers and gateways providing system of record capabilities. In addition to data persistence, other data-centric activities include data enrichment and data analysis. The analysis sets the stage for decisioning and business processing that support actions. Finally, closed-loop processing allows the opportunity for learning and business process improvement. It is possible for back-end systems to interface directly with the device tier, but this will only be the case for applications that do not have high scalability requirements. As we already noted, some or most of these activities can also be performed by the middle tier.

# Figure 3: IoT edge computing model with gateways

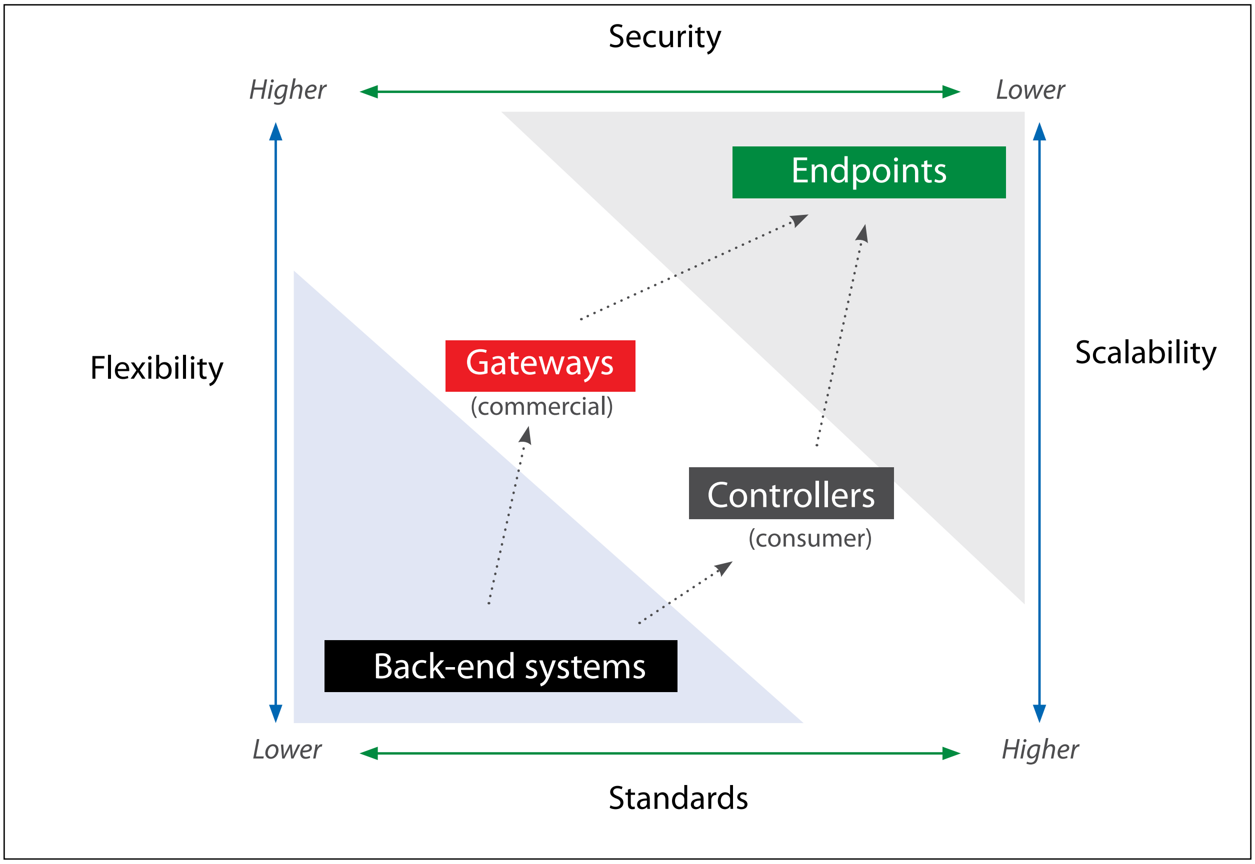


An architecture above, developed by Stephen Hendrick at Enterprise Strategy Group (ESG), describes a three-tier model that provides high level of flexibility and is already finding its way into a variety of commercial and consumer products. Scalability, security, standards, and flexibility are key design points of this model.

The need for a middle tier is contingent to the style of data collection that occurs in the front-end. Some sensors "continuously" send data while others only send data when specific conditions are met such as exceeding a threshold. While continuous data feeds provide a richer opportunity for understanding trends and patterns, a great deal rides on the specific use case. For example, a moisture sensor under your HVAC system doesn't need to continuously report data—it just needs to send an alert and shut off the system if it is triggered.

Alternatively, the Fitbit on your wrist is designed to "continuously" collect data so that back-end analytics can help you understand daily and monthly trends across a variety of biometrics.

Middle-tier gateways and controllers are needed to manage high volumes of continuous data. Network efficiency, security, and workload distribution are the primary drivers for gateways. High expected IoT data volumes means efficient use of networks is critical. The best way to reduce network traffic is to send less data, so compression is a logical starting point. But gateways can be even more effective at reducing traffic by acting as local drop boxes for collecting, aggregating, and summarizing data before shipping the results to the back-end. Gateways also can provide an added layer of security and isolation for back-end systems. Depending upon system design, the gateway can offer the ability to offload and distribute some back-end processes, especially those associated with filtering and analytics. Centralization, consolidation of data, and process coordination are the primary drivers for using controllers in a consumer environment.

A controller will often be located in the home and combine data collection, analysis, and processes the drive actions. So the controller will collect data specific to presence (biometrics) and the environment (natural and physical) and leverage user-defined policy to make decisions and recommendations while continuously adapting to changing preferences and environmental factors. Consequently, the existence of a middle tier with gateways and/or controllers helps enterprises deliver IoT experiences that are secure, scalable, flexible, and adhere to standards.


# Edge analytics for asset performance and predictive maintenance

**Name of Organization:** Envision
**Industry:** Energy
**Location:** New York, NY USA

## Business Opportunity or Challenge:

Envision Energy, one of the world's ten largest wind turbine companies, has a growing Big Data challenge. Along with focusing on producing wind turbines, the company also has a software business which manages up to 13 gigawatts of renewable energy assets—both wind and solar—globally.

The energy business can no longer be differentiated just by mechanical engineering, but by an ability to monitor and maintain high performance of energy assets. Each of Envision's 20,000 wind turbines are built with over 150 advanced sensors that continually assess acceleration, temperature and vibration of the turbines. Extracting data from their wind turbine sensors lets them see trends and create predictions for performance optimization and maintenance to minimize downtime.

Envision's Smart Wind Farm System is an integrated asset management solution, providing predictive and preventive services and maintenance support to optimize performance.

**ENVSION MOVED FROM ANALYING DATA EVERY 10 MINUTES TO EVERY FEW SECONDS.**

Strategically developed from the ground up by wind experts to incorporate industry best practices and cater to the unique needs of wind farm operators.

"Envision is a renewable energy company in the wind-power business, but we also have a software business," says Dr. Guido Jouret, President of Envision Digital Innovation Center, in a video. "So we're becoming a smart-energy company. "Our main business challenges are that we operate devices like wind turbines and solar panels, which are very complex, and we have to operate those devices for decades—which means maximizing availability, but also maximizing power."

With more than 3 million sensors, the company is managing massive amounts of real-time data—more than 20 terabytes of data at a time—to help continuously monitor this vast wind turbine network. In addition, the data volume is growing at over 50 percent annually as Envision continues to collect more data, more frequently from each of their wind turbines. To complicate things even further, the wind farms which house their turbines are geographically dispersed.

## How This Business Opportunity or Challenge Was Met:

The key to managing all this real-time data and making it available to the business boiled down to analyzing turbine sensor data with greater granularity, Envision's technology planners concluded. To better address this challenge, they have moved from analyzing turbine data every 10 minutes to every minute, and then on to every few seconds. By immediately analyzing real-time sensor data from their wind turbines, Envision is able to quickly identify actionable insights with significant business benefits.

"We're taking the raw data and turn it into trends," says Jouret. "Then we can take the trends and turn them into predictions about what will happen based on what we've seen in the past, to minimize downtime and take advantage of performance increases."

Envision needed an analytics solution which enabled them to handle these multiple terabytes of data with sub-second response time, along with the capability to run distributed queries/edge analytics closer to the source of data. The company also needed to be able to continuously import and store large amounts of real-time sensor data with the ability run fast and flexible queries locally, in their central data center, or in the cloud. "We wanted to run the data queries in our data center or in the cloud," Jouret says, adding that Envision's goal was to employ analytics to help "reduce our customers' total energy output by 50%."

The energy company employed ParStream's Analytics Platform, which handles massive volumes and high velocity of IoT data, to manage data streaming from its highly distributed network. "Our wind farms and solar farms are all spread out," says Jouret. "We needed to ingest the data at the source and also run queries against huge data sets."

**EDGE ANALYTICS DELIVERED AN OVERALL 15% IMPROVEMENT IN PRODUCTIVITY FOR A WIND AND SOLAR NETWORK.**

Envision continuously improves the mechanical aspects of its turbines, "but first examines historical operational data for potential design improvements," according to Syed Hoda, chief marketing officer for Parstream. "Envision's strategies around 'how many sensors?', 'where to place sensors?', 'what data to collect and how often?' are perfect examples of collaboration between the physical and digital worlds."

## Measurable/Quantifiable and "Soft" Benefits from This Initiative:

As a result of employing the distributed analytics platform, Envision was able to address both performance optimization and predictive maintenance within its wind and solar network, which combined, help to deliver an overall 15 percent improvement in productivity.

For performance optimization, Envision uses sensor data to make smart decisions about altering the angle and speed of the turbine blades in order to optimize performance at any given time, based on changing environmental conditions. "Through the use of real-time sensor data, we can boost a customer's total energy output by up to 15% from their wind farms," says Jouret.

In addition, for predictive maintenance, Envision's sensor technology helps the company perform checks for any irregularities in operational performance for its 20,000 wind turbines, enabling technicians to predict potential failures before they happen. Real-time data is matched against historical data to determine which parts need adjustments or replacements, significantly reducing downtime.

Case study

# How do you remotely service ATMs?

**Name of Organization:** Diebold, Inc.
**Industry:** Technology
**Location:** Canton, OH USA


## Business Opportunity or Challenge:

Automated teller machines are everywhere and ATM repair is a serious challenge, as machines can frequently go out of service, meaning lost revenue for the ATM operator.

Diebold, Inc. is a global leader in providing integrated self-service delivery and security systems and services, with ATMs in almost 90 countries. The vendor's service channel accounts for more than half of its revenues. With a great deal of pressure on the financial services sector to reduce operating systems, Diebold needed to rethink its standard service-and-repair model, which was onsite.

As recounted in case study published by PTC, Diebold sought to devise a way to remotely service, in real time, its increasingly software-driven self-service ATMs while adhering to iron-clad security practices. The goal was to service ATMs and respond to issues as they occur, versus dispatching technicians to travel to the physical site to address the problem.

## How This Business Opportunity or Challenge Was Met:

To enable real-time remote servicing of the ATMs in its far-flung network, Diebold adopted a remote service model for its Opteva ATMs known as OpteView. The ATMs employ PTC Axeda software to create a service approach that is connected, allowing a direct communication flow between Diebold and end-user ATMs.

"When we released our Opteva platform, we designed in the ability to remotely service and diagnose our product because our traditional service model was based 100 percent on the dispatch of parts and labor to the site whenever the ATM failed," says Paul Mercina, director of service product management/operations planning at Diebold. "With this remote service technology, we can start diagnosing problems at the time of failure, and, in some

**OPTEVA ATMS HAVE BUILT-IN DATA-CAPTURE TECHNOLOGY AND STORE PERTINENT INFORMATION ABOUT A DEVICE'S PERFORMANCE FOR REMOTE ACCESS.**

cases, actually correct the failure without waiting for a technician to go on site, which could be within minutes."

Opteva ATMs have built-in data-capture technology and store pertinent information about a device's performance for access by OpteView. Diebold's remote service support operators are directly alerted when a problem with any ATM occurs.

OpteView enables regularly scheduled data collection points. This data is compiled and reviewed periodically to enable Diebold to accurately track device performance metrics. Then, predictive rules and processes can be established to provide the customer and the manufacturer with advance notification of which components will fail and at what time.

## Measurable/Quantifiable and "Soft" Benefits from This Initiative:

Diebold's traditional service model does not begin until a customer service engineer (CSE) arrives on site. For ATM repair and servicing, Diebold can now remotely access the ATM either while the CSE is en route to the ATM or before the CSE is dispatched.

OpteView remote support operators can conduct a remote diagnostic session to provide detailed information to the CSE when needed, which reduces the onsite resolution time while increasing the probability that the CSE will remedy the issue the first time and reduce repeat site visits. Ideally, the problem can be corrected remotely, entirely eliminating the need for a CSE to visit the site.

With real-time ATM data, service staff can diagnose and begin resolving problems immediately. As a result, the level of information on the machine's status is much more detailed and the ATM can be analyzed before a technician arrives on site. This enables Diebold to deploy technicians with precise knowledge about individual machines, increasing "first-time-fix" ratios.

With the remote diagnostic and repair capabilities of OpteView, Diebold can expand the level of its service and support offerings to include predictive maintenance, software version control, and remote monitoring and notification. So far, the PTC case study reports, about 17 percent of problems have been resolved remotely, with a 15 percent decrease in downtime. Turnaround time for problem resolution has been reduced from roughly three hours to fewer than 30 minutes when remote corrective action is successful. When a CSE must be deployed, response time is guaranteed, in accordance with the customer's service contract.

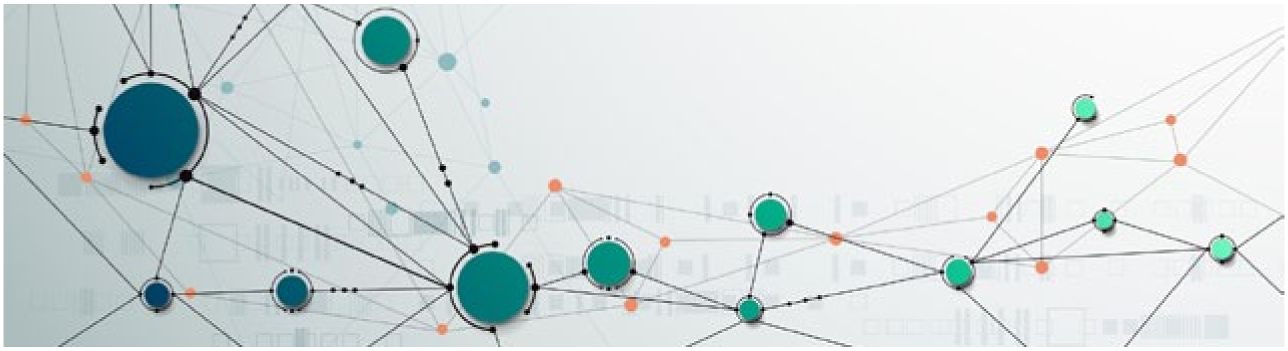

RTInsights.com

*Accelerate Your Business with Real-Time Insights*

RTInsights is an independent, expert-driven web resource for senior business and IT enterprise professionals in vertical industries. We help our readers understand how they can transform their business to higher value outcomes and new business models with IoT real time analytics. We provide clarity and direction amid the often confusing array of approaches and vendor solutions. We provide our partners with a unique combination of services and deep domain expertise to improve their product marketing, lead generation, and thought leadership activity.