influxdata®

AN INFLUXDATA CASE STUDY

How eSmart Systems Uses MS Azure and InfluxDB Enterprise to Optimize Energy Investments

Eric Asberg
CTO, eSmart System

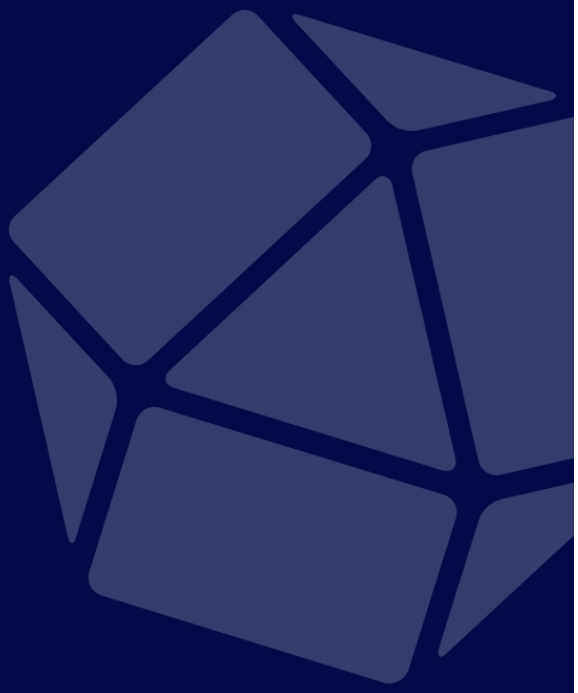

eSmart SYSTEMS

JUNE 2018

Company in brief

eSmart Systems develops digital intelligence for the energy industry and smart communities. It is based on more than 20 years of international experience in establishing and operating knowledge-based, leading IT and energy related companies targeting global markets.

Founded in 2012, the company has offices in Norway Denmark, the UK, and USA. eSmart Systems provides software solutions to the energy industry, service providers and smart cities. Its platform is designed to handle and exploit the Internet of Things, Big Data and Analytics in real time. It has broad application possibilities that give grid operators insight into their distribution network, make energy trade between local market prosumers possible, and monitoring city air quality easy.

Common to all applications is vast data quantities gathered from sensors, which are analyzed using advanced prediction and optimization models. This results in completely new ways of visualizing data, making decisions and saving resources and costs.

eSmart's mission is to build digital intelligence to provide exceptional solutions to its customers and accelerate the transition to sustainable societies.

Case overview

eSmart Systems wanted to put Big Data and IoT technologies to work for its energy customers to optimize their investments and enable next-generation operational performance. With industry changes such as smart meter and renewable energy adoption, utilities companies needed to make data-driven decisions to improve efficiency and cut cost.

eSmart Systems use MS Azure and InfluxDB Enterprise to gather vast amounts of data from sensors and analyze it using advanced prediction and optimization models. This results in a completely new way of visualizing data, while helping their customers make decisions faster to save resources and costs. By running InfluxDB at scale on Azure using InfluxDB Enterprise to meet data privacy requirements, eSmart Systems is converting data into asset management decisions for its customers.

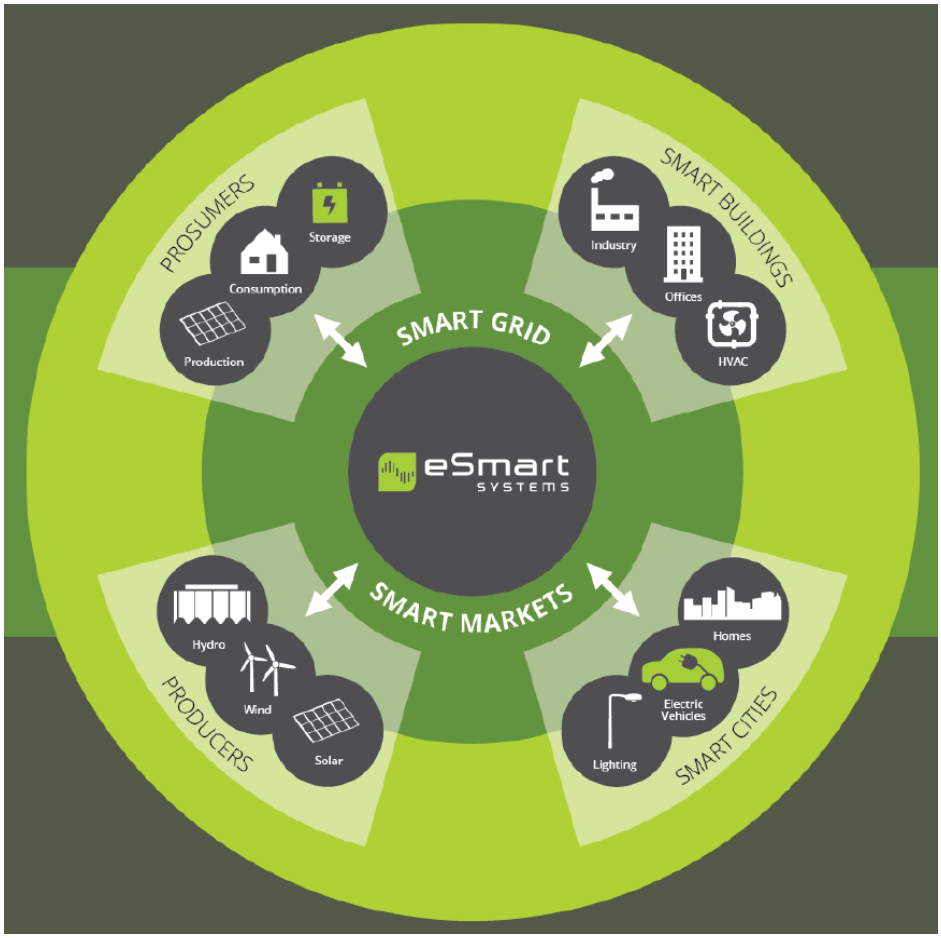

InfluxDB powering eSmart Systems' multi-segment platform

> *"We wanted to use this new technology to solve these new energy industry challenges."*
>
> ***Erik Åsberg,*** *CTO*

The business problem

The energy industry is undergoing changes such as new types of power-demanding loads like Electric Vehicles (EVs) and the increasing popularity of distributed generation like Plug-in Electric Vehicle (PEV) and wind generation. Meanwhile, the IT industry is also undergoing massive change, witnessing the rise of new technology to handle Big Data, Artificial Intelligence, and Machine Learning. eSmart wanted to use these new technologies to solve these new energy industry challenges and thereby offer its customers actionable digital intelligence in real time.

Challenges facing utilities providers include an aging power infrastructure, very long power lines, and difficult weather conditions in many parts of the world. This is compounded by the inefficiency of traditional power line inspection methods:

- Current inspection methods are slow, dangerous and manual.
- When data is gathered, it is being manually processed which is very time-consuming.
- Fixing power outages as quickly as possible is absolutely vital for utilities (not only to maintain power supply but also because utilities in countries like Norway are penalized if they don't deliver power to end users, so outages are a double loss for them).

To solve these problems, eSmart turned to AI, Big Data analytics and drones:

- Using Deep Learning to find problems automatically
- Using drones as "the eye in the the sky"
- Developing a tool for field crew that makes their job easier and safer
- Attempting to predict problems before they turn into critical errors

They set out to achieve three objectives:

1. Reduce the need for investments in the power grid by increasing the utilization factor of the existing infrastructure
2. Provide high-quality predictions and analytics, enabling their customers to operate more efficiently and remain one step ahead of the competition
3. Make use of IoT and Big Data technologies to enable predictions

eSmart Systems Objectives

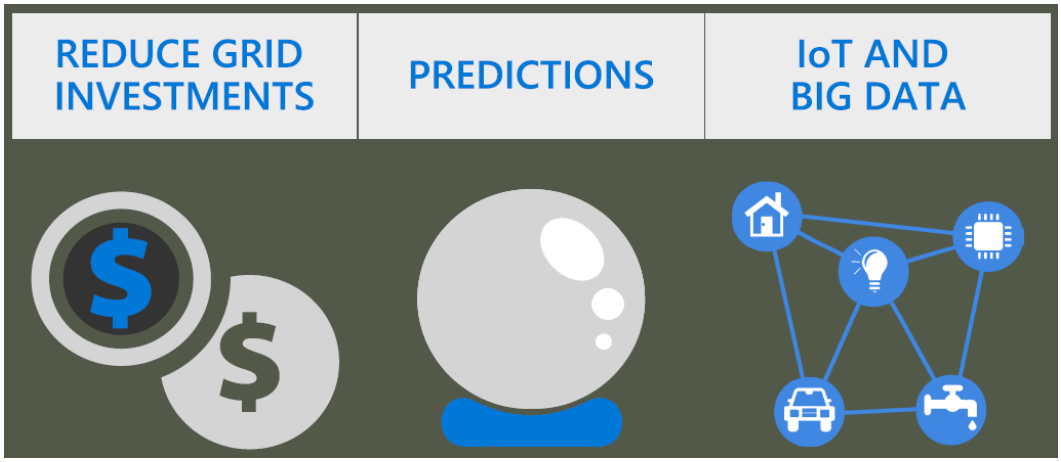

To achieve these objectives, eSmart Systems apply a three-prong digital intelligence process aimed at delivering value to customers where the three prongs intersect.

eSmart Systems Digital Intelligence Process

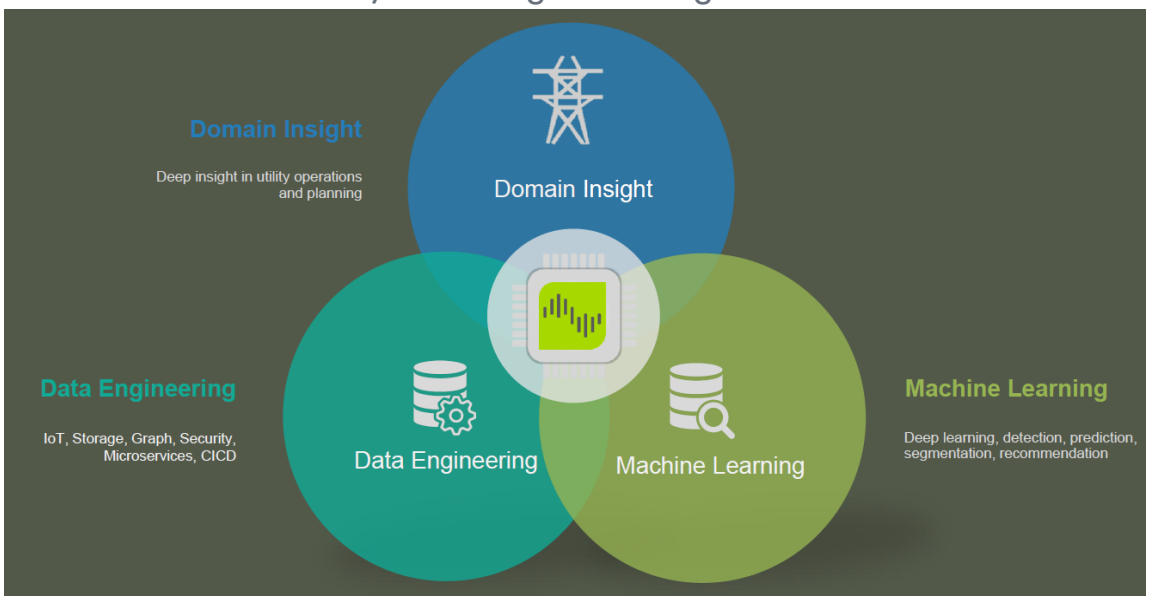

Within the field of Machine Learning and Big Data analytics, they use a variety of methods but focus on the latest in deep learning and work very closely with Microsoft expert teams.

How eSmart Systems Use Machine Learning & Big Data Analytics

Load predictions	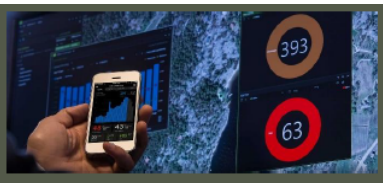	• Substation load • EV Charging • Peak loads after outage
Segmentation & profiling	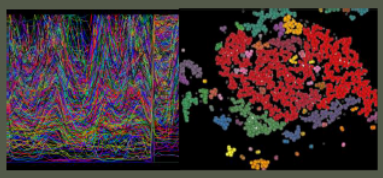	• Customer behaviour • Identify household that install solar and/or acquire EV
Risk monitoring		• Aggregate data • Estimate risk for outage • Estimate meter failure
Failure & anomaly detection		• Identify components based on object recognition tech. • Identify failures & anomalies

When eSmart Systems' AI team started looking at their predictions at scale, they faced challenges with their own existing solution (based on Azure Blob Storage, which they had evaluated 5 years ago). They decided to re-evaluate it in light of today's technologies.

The technical problem

eSmart Systems have one data platform to support all six products in their three business areas.

One Platform for Three Business Areas

eSmart Systems' Six Products by Business Area

The company's founding team had worked with energy and IT for over 20 years. Working in a traditional energy setup with a relational database behind a Windows frontend, they witnessed the volume of data growing and experienced the endless cycle of tuning and re-tuning the databases as data grew, and accordingly, the constant need to buy new hardware. They saw that setup's lack of scalability. So when they had the chance to start fresh and leave their legacy systems behind, they knew they had to go

cloud-based — and build a solution designed for utilizing the flexibility and elasticity of the cloud. They chose Microsoft Azure as their primary platform because they felt that it was more enterprise-ready than competitors.

Designed for Big Data from the ground up, the eSmart Systems platform is a generic platform designed to handle time series in real time and able to support a variety of business opportunities. Utilities are their major business area, yet as energy industry challenges evolved, eSmart Systems had to upgrade their platform to provide predictions and digital intelligence at scale in real time and realized that they needed a time series database.

The solution

> *"We need to make sure that the database is always backed up, that it's always patched up, that it's always up and running. We don't want to spend our time doing that stuff because really, we need to use all our time in making value for our customers on top of those services."*

Why InfluxDB?

Through the AI work that they did, it became quickly obvious to eSmart Systems that InfluxDB was the time series database they needed:

- Time series data (since it shows the rate of change of a given parameter over time) is central to machine learning and AI, and InfluxDB is a purpose-built time series database.
- To start training your algorithms, you need a dataset where the answer is contained. A lot of the data that they wanted to use was in the form of time series — in the form of time-stamped data collected from sensors & smart meters as well as weather and calendar information.
- Data availability was the key determining factor for choosing InfluxDB. InfluxDB running at scale made it very easy for them to make good predictions with the type of competence that they have amongst their data scientists.

The Connected Grid

eSmart's first and most mature product, The Connected Grid has the following functionality:

- Single pane of glass into the various systems that collect data about their customers operations, including infrastructure and their customers
- IoT data management (IoT devices, smart meters and other types of sensors in the power grid, and on the quality of that data)
- IoT data collection (eSmart provide high-level data collection dashboards for their customers, displaying KPIs and statistics and showing which work orders failed to deliver target quality). For eSmart, capturing the data is the starting point for everything:

    1. New types and latest generations of smart meters present new problems for utilities, such as data usage methods and cost and privacy.
    2. It's tempting to examine every data point that a smart meter is able to produce, yet this often involves a mobile connection and incurs unnecessary costs.
    3. eSmart help their utilities customers get an overview of their data usage.

Connected Grid - Data Usage Overview

Often in the utilities use cases, the devices that you need to collect data from are in remote areas, requiring a mobile communication channel to send and receive data. And since each mobile connection comes with a monthly charge, it is important to also track the data usage across these connections. In the upper left corner of the above image are SIM cards that have exceeded their data plan, and in the lower right corner are those in good standing. eSmart help customers either identify which SIM cards should undergo a data plan change or revisit how much data they collect for that specific case.

Overall, the the Connected Grid provides customers with:

- Robust load predictions (system advice and actions for operators to take)

- Aggregations, validation, estimation, and editing of values
- Flexible workflow editor (every operation within the system is defined by an editable workflow)
- Events and alarms that eSmart collect from IoT devices in real time

Connected Grid - IoT Data Collection

- **Transformer load management** (In cold climates where electricity is used for heating, transformers tend to have load challenges, and when overloaded, pose the risk of an outage — or worse — of fire or explosions). To help utilities companies avoid the safety hazard of overloaded transformers, they:

    1. Collect data from the smart meters connected to grid and data from the transformer and from the sensors on the transformer itself; create a virtual sensor by aggregating smart meter data; and compare those aggregations to the actual metering done on the transformer itself
    2. Help customers understand which transformers are overloaded and which have free capacity, and thereby help them get more out of their infrastructure
    3. Deliver digital intelligence that is enabling many of their customers to optimize transformer usage

Connected Grid - Transformer Load Management

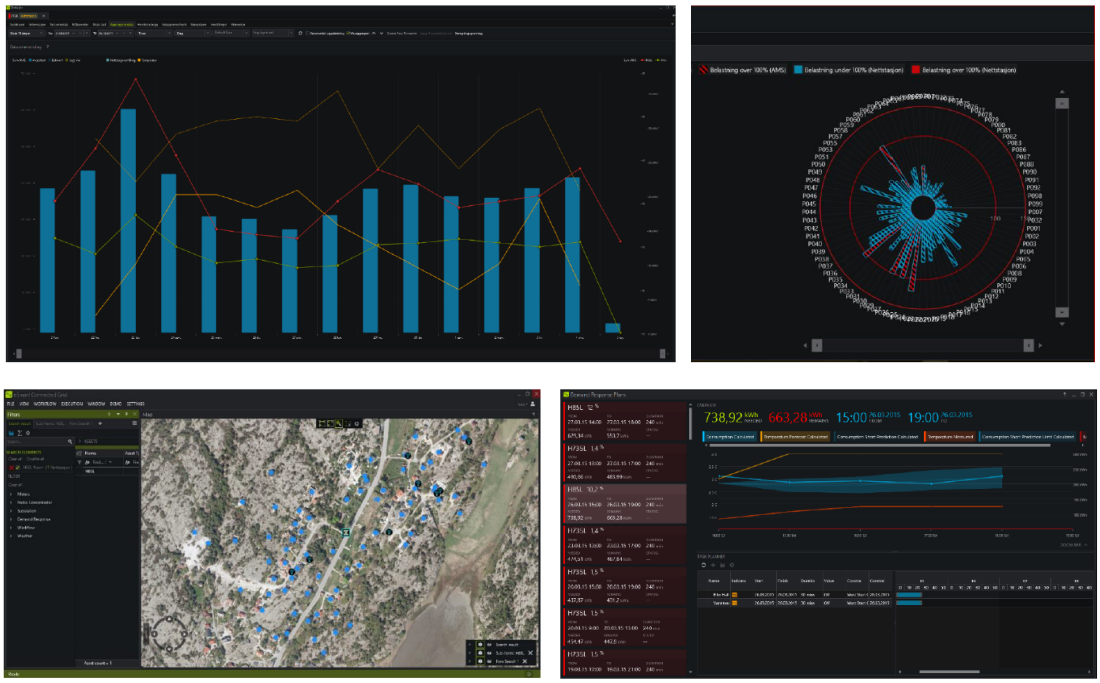

As an extension of this, eSmart can also perform control signals. They have connections to smart homes, and home gateways that they can control. They gather real-time information from usage in them. They then use their predictions to determine whether it's likely that a given transformer will be overloaded. If they identify an overload risk, the system will calculate how much power is needed to shut down to prevent that overload and then execute that control plan.

- **Work order management:** eSmart perform work order management from A to D (rather than A to Z) because customers' work order system is often one that their field operators have long been using and that is not easily replaceable. Instead eSmart integrate with customers' existing work order management to provide an overview of work order status.

The Connected Drone

Since the Connected Grid holds information about every asset and component in the grid, as an extension to that, eSmart Systems decided to get real-time information from the use of drones to provide accurate information about each component's status. They take the image captured in the field and do the analytics, map them with the correct component in their component registry as well as give customers an overview of all their inspection findings. It's much easier to use a drone than the standard inspection method of using helicopters, which are weather-dependent and much less accessible.

Connected Drone

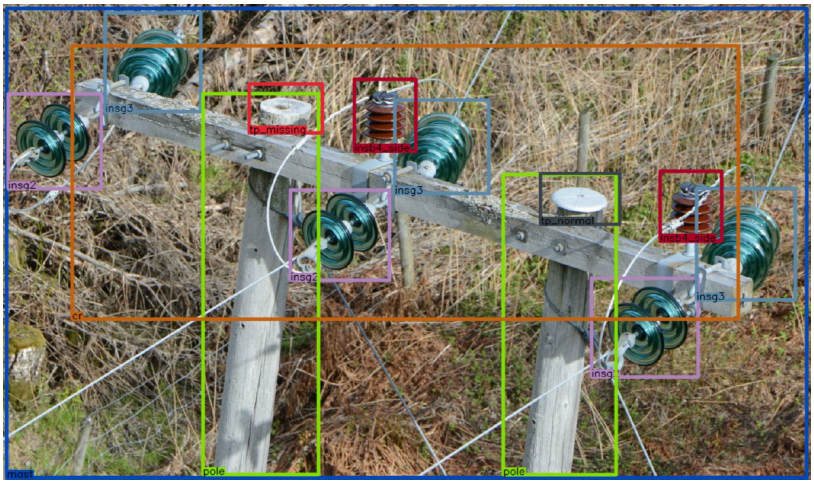

Though utilities providers have numerous pictures of their assets, pictures of anomalies are not so common. Yet that is what machine learning algorithms need to identify anomalies. Therefore eSmart create these anomalies themselves by taking real components, digitizing them, and putting them into a gaming engine to produce synthetic data in order to train their AI. Generative Adversarial Networks (GANs) are used as image refiners to make synthetic images more realistic and more suitable for training, thereby enabling anomaly identification without having many real pictures.

GANS for Synthetic Images

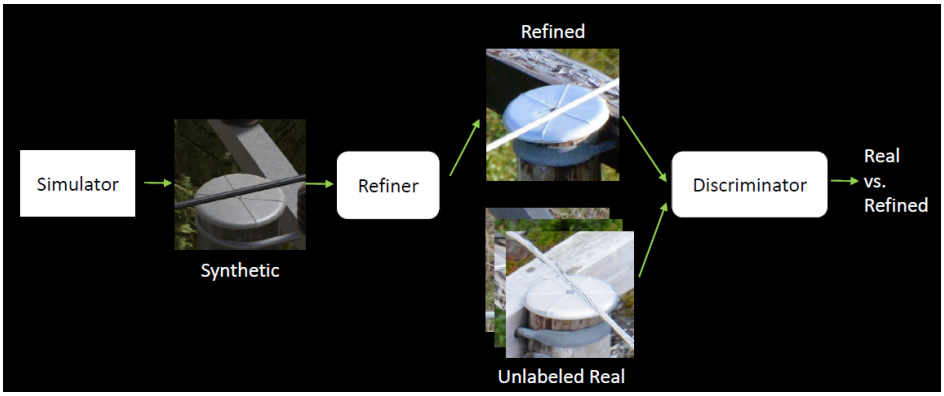

Technical architecture

> *"We do substation load predictions, EV charging predictions, and predict peak loads after outages. It was through this work, really, that InfluxDB came up as a source for us for storing and utilizing time series better than our previous solution."*

Data Delivering a SaaS Using InfluxDB

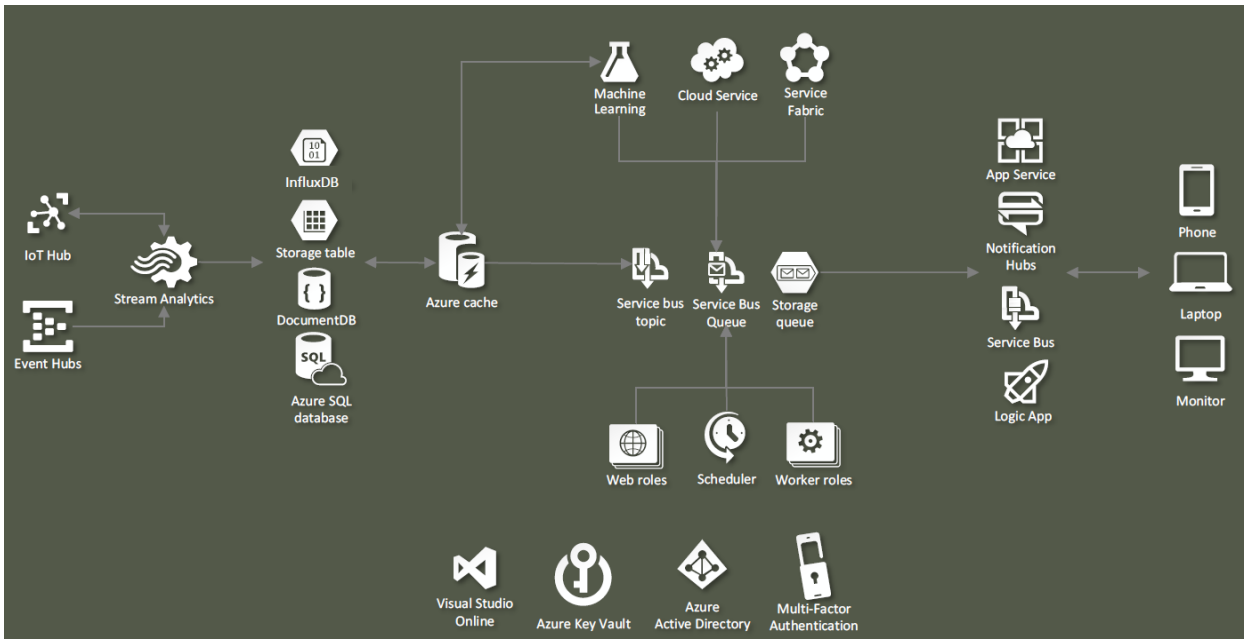

eSmart Systems run InfluxDB at scale on Azure by using InfluxDB Enterprise. The sensitive data they collect is from smart meters, showing power usage within a private household. Energy usage can tell a story about a person or their household — such as whether the person is at home or how much they're consuming.

eSmart Systems created the concept of a top system. On the lower side is a typical ecosystem within the utility. Utilities have a lot of systems that are extremely good at what they do but are siloed, so they focus only on their task. So when work processes start to span across these silos, it becomes difficult to utilize all the data available because context from other systems is needed. That's where eSmart Systems come in. They ingest data from all these lines of business systems, as well as directly from smart meters, sensors, drones, and external sources like weather information, even monitor social media on behalf of their utilities, and market information — to enrich available data and get more value out of it.

They use their data model to find correlations and perform analytics on top of existing data.
Their top system concept (based on four principles detailed below), is a cloud above their existing infrastructure — a layer between their line of business systems and their IoT devices.

Top System Concept

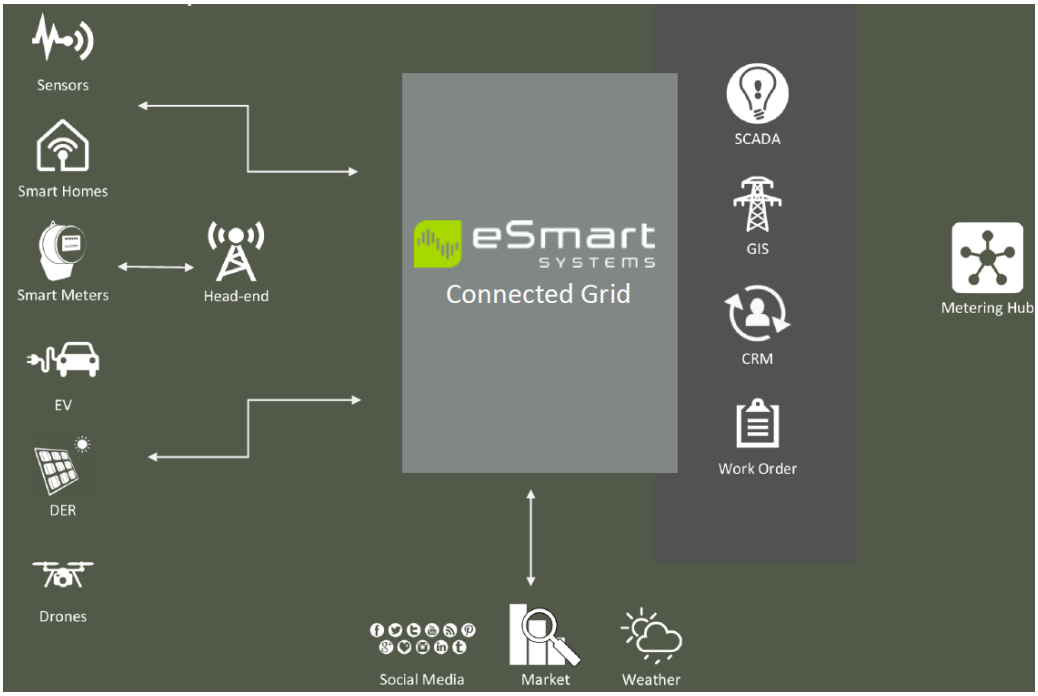

The eSmart Systems top layer system utilizes internal data sources enriched with external data, together with Artificial Intelligence, to extract new insights for operations and executives.

- They store all time series in InfluxDB. They still store the image in Azure blob storage but with a reference into InfluxDB (a timestamp and a value).
- They can then use those two sets of data to understand if there was a change in the actual thing that they're looking at versus the image.
- They can then compare what happened around that time of the image or of that tweet. (certain values or certain alarms triggered at the same time) to establish that timeline, and InfluxDB is core in that timeline.

The platform is built on a PaaS solution:

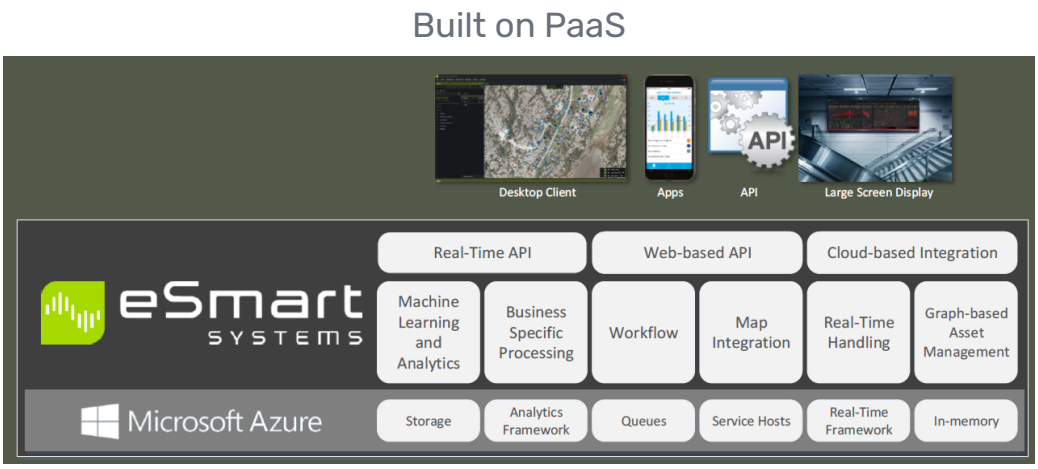

eSmart Systems developed their platform's architecture based on four principles:

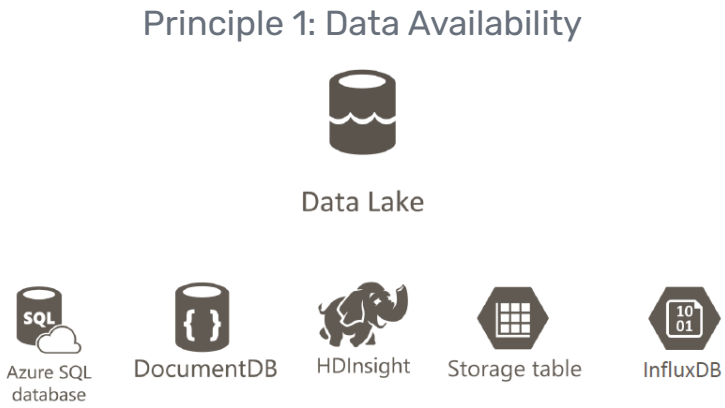

eSmart Systems think about the data structure within their Data Lake and utilize all types of technology to make their data available — ensuring that they choose the right technology for the right type of data. That became possible for time series data as well since, with InfluxDB, storage was no longer a cost issue.

Principle 2: Time Series

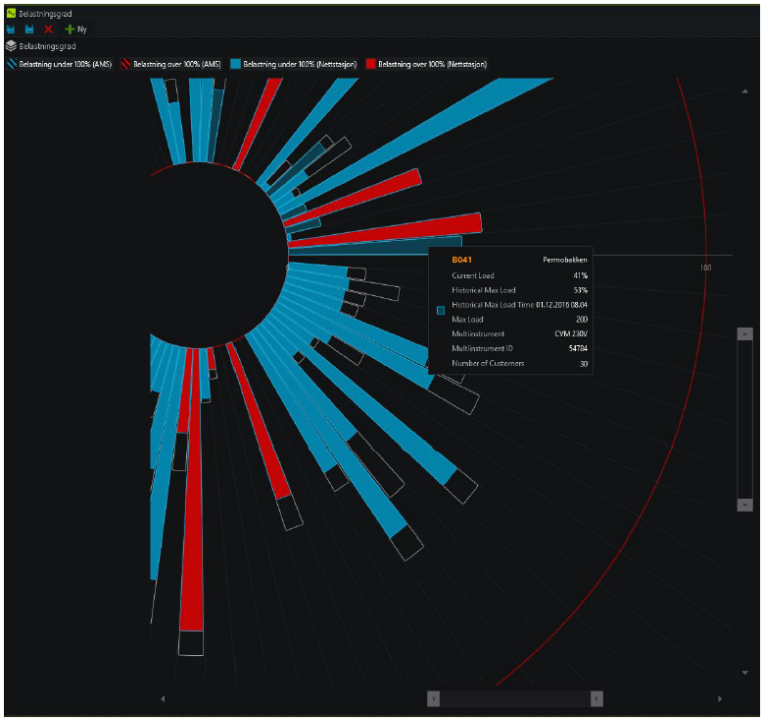

The second principle is time series. eSmart Systems have an extreme, wide definition of time series. Of course, that includes metering values from sensors. But they also consider — as time series — the alarms coming from IoT devices, messages from customers, tweets, and images from drones because they want to establish a timeline. Machine learning and analytics help them better understand what is about to happen now because they have established a timeline with this extreme definition of time series.

Principle 3: Graphs

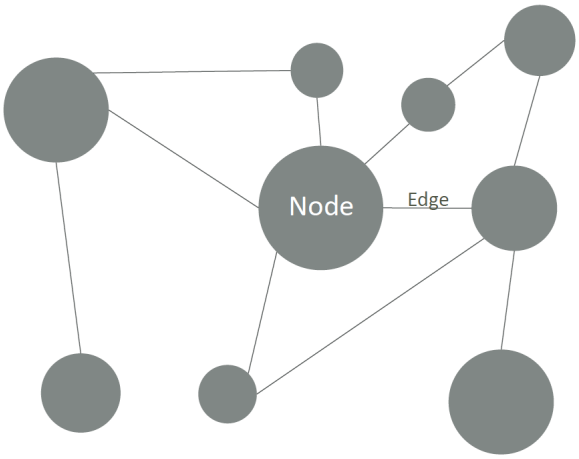

The third principle is user graphs. Graphs provide an overview of the physical infrastructure. Since IoT devices do not necessarily communicate in the same way, they have another layer in the graph that tracks communication paths between devices. Multiple graph layers help establish the timeline as well as learn the connections between all the assets.

Principle 4: Integrations

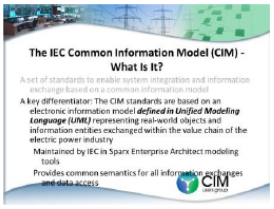

The fourth principle is integrations. For eSmart Systems, it's very important to avoid data lockups while also getting hold of the data. They use the energy industry's Common Information Model (CIM), which makes it easier to share data between their systems.

What's next for eSmart Systems?

eSmart Systems are looking at how InfluxDB Enterprise can further fit into their architecture and how to best utilize Telegraf and Kapacitor.

Results

> *"By just picking up the phone, spending two hours with the tech guys at InfluxData, we learn so much more than trying to understand the details ourselves."*

eSmart Systems' intelligent analytics platform captures, analyzes, visualizes and converts real-time operational data into actionable insights to enable next-generation operational performance. Connected Grid gives customers all the information in a single system, which significantly improves their work processes. It is saving customers millions by using predictive AI on sensor data. For example,

eSmart System customer Jacksonville Electric Authority (Florida) had to send a truck out and dig up the water meter to check its condition when a smart water meter has provided zero values for over 3 months. Yet in many cases, the water meter turned out to be in perfect condition. So the customer wanted a system that can reduce the number of truck rolls because they are very costly. eSmart reuse smart electricity data together with smart water data and then use machine learning to predict likelihood of the water meter being broken. Jacksonville Electric Authority have reduced their truck rolls by over 80% and saved over $700,000 per year just by this simple comparison of data. eSmart use the same approach with district heating within Connected Grids because it has many similarities in the infrastructure. Connected Drone is a true game changer for power line inspections. It utilizes Digital Intelligence applications such as image recognition to drastically improve infrastructure inspections. With Connected Drone, utilities for the first time have the ability to use Artificial Intelligence to catalog infrastructure components through eSmart's Intelligent Assistant. Running on eSmart Systems Connected Platform and Microsoft Azure, The Intelligent Assistant can analyze 100,000 images in less than an hour — that's more than a human can do in a year!

Today eSmart Systems move fast, act agile and grow with their customers. With a SaaS-based business model and value-proven solutions, they save customers time and money. Using InfluxDB Enterprise on Azure, eSmart Systems is fulfilling its vision as the "provider of next-generation IT solutions".

About InfluxData

InfluxData is the creator of InfluxDB, the leading time series platform. We empower developers and organizations, such as Cisco, IBM, Lego, Siemens, and Tesla, to build transformative IoT, analytics and monitoring applications. Our technology is purpose-built to handle the massive volumes of time-stamped data produced by sensors, applications and computer infrastructure. Easy to start and scale, InfluxDB gives developers time to focus on the features and functionalities that give their apps a competitive edge. InfluxData is headquartered in San Francisco, with a workforce distributed throughout the U.S. and across Europe. For more information, visit influxdata.com and follow us @InfluxDB.