influxdata®

AN INFLUXDATA CASE STUDY

How Factry's Industrial Historian Solution Helps A&S Energie Generate Efficiently and with More Confidence

Frederik Van Leeckwyck

Business Development Manager, Factry

factry

JUNE 2018

Overview

Factry wanted to help their customer A&S Energie achieve greater efficiencies with energy generation by using a modern process historian (historians are software used to collect process data from production sites). A&S Energie's existing historian was underutilized, and data migration and collector setups in a new historian solution would have to be done securely in line with plant network security policies. Factry replaced A&S Energie's existing solution with Factry Historian, powered by open source InfluxDB. This allowed A&S Energie to increase use of the process data to make their staff more efficient as well as find failures faster.

Factry Historian collects and processes data from A&S Energie's industrial system, stores it in InfluxDB, and visualizes it in Grafana. Migration to Factry Historian resulted in enhanced performance visibility and plant management for A&S Energie.

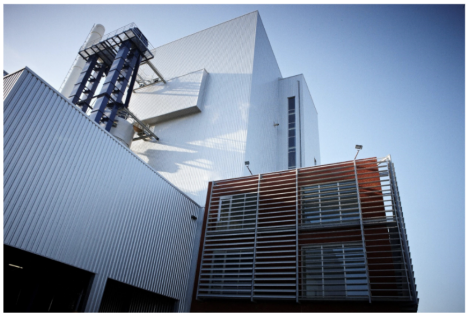

Increased visibility and efficiency gained by A&S Energie upon adopting Factry Historian powered by open source InfluxDB

About Factry

Factry is a Belgian software company that helps its customers gain insights into their operations by bringing the power of open source software to the world of Industrial IoT. To make this possible, Factry built Factry Historian, a solution which collects process data from industrial systems and stores it in InfluxDB.

Factry was founded by ex-Siemens Manufacturing Execution System specialists who realized that they could build a better historian using open source technologies. Factry is focused on the process industry (both real physical products, such as food processing, and intangible processes that run continuously, such as energy generation).

Factry's mission is to achieve data-driven operational improvement for its customers. It offers an open platform for easily collecting, storing and visualizing sensor data from industrial equipment. Together with its transparent approach, this guarantees modern and performant solutions without artificial limits.

> *We benchmarked several Time Series Databases and quickly found that InfluxDB was the most powerful for our use case.*

Frederik Van Leeckwyck, Business Development Manager

The business problem

A&S Energie is a Belgian biomass plant that supplies electricity to 55,000 households by burning 180,000 metric tons of non-recyclable wood per year—that's 22 tons per hour. At its maximum output, the plant generates 26 megawatts of power. The plant has been operational since 2010 and is operated by a team of 25 comprising management and operators.

A&S Energie presented a classic Industrial IoT challenge involving high-cost, confined industrial equipment requiring a useful life of many decades. A&S Energie consists of a single plant. At its center is a huge "kettle", nine stories high and constantly fed with non-recyclable wood that is burned to generate steam. Steam passes over a turbine and generates electricity. Since all measurements and sensors are present in the plant, data collection is also done locally.

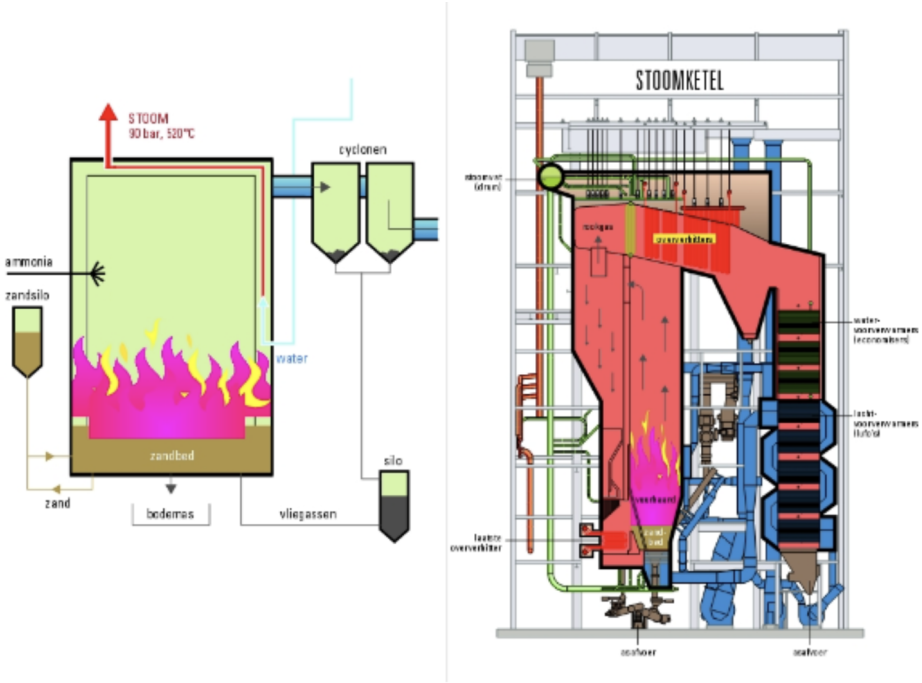

Prior to Factry's intervention, A&S Energie relied on a largely manual, error-prone, and time-consuming process:

- Operators wrote down counters on paper (values of certain measurements needed for reporting on the power plant's status).
- Based on this paper-reliant reporting, Management reporting was done with Excel.
- The existing Historian was underutilized—its data was not being used to extract meaning.

Factry was hired by A&S Energie to help management optimize operations and enable operators to work with process data more fluently. In essence, this required:

- Drastically simplifying reporting on critical parameters for maintenance and operations
- Increasing the usability of process data visualization

Factry first conducted a review of A&S Energie's existing historian solution; a solution that came as part of the service contract that they had with the original equipment that they purchased. Factry's review concluded that the existing historian was used only to passively store data and offered no further insights from the data. Further, the historical view only provided 2 weeks' worth of data, which isn't enough data to identify how A&S Energie could streamline operations.

The technical problem

Power plants, like manufacturing companies, are typically controlled with the help of Supervisory Control And Data Acquisition (SCADA) systems. These systems give operators a holistic view of the plant and are used to control and monitor the equipment in use at the plant. Operators use SCADA systems on a daily basis in the control room.

SCADA systems are typically supplemented with a process historian, which is used to archive sensor values, setpoints, and alarms for later analysis and real-time trending. For A&S Energie, their systems produce about 2,000 relevant data points at 1 second resolution kept in long-term storage.

Typically, this type of data is kept for multiple years in its original form without any downsampling. This is necessary to provide sufficient levels of detail when troubleshooting incidents, even years later. For example, a 3-second increase in pressure due to a failing valve might get overlooked when the data is averaged at minute intervals.

Being focused on the process industry, Factry had selected early on which technologies to use in their products. To gain insight into how processes are running in a factory or energy plant, they realized that two types of data are needed:

- Time-based data - such as time-stamped sensor values
- Relational data - such as KPIs for use in management dashboards, etc.

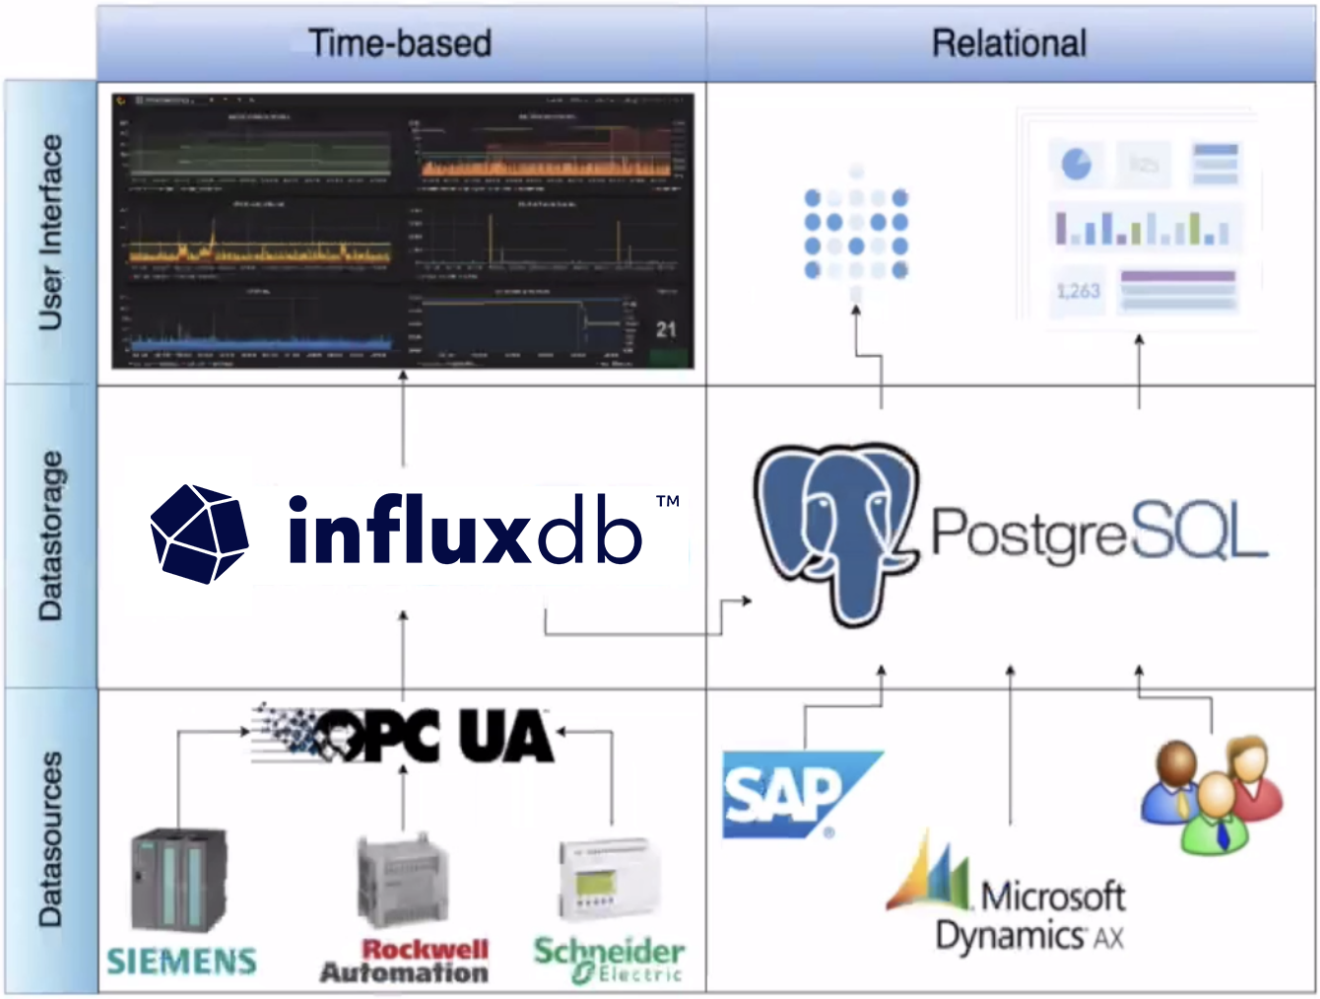

So they decided to deploy the technologies illustrated above in building Factry Historian. For storing time-stamped data, they benchmarked several Time Series Databases, and chose open source InfluxDB to be Factry Historian's storage backend.

The solution

> *Factry Historian provides real-time and historical insights for everyone, from machine operator to plant manager. The technology is really advanced. It does amazing things. And it has become so powerful that everyone can start using this and gain insights into how processes are running.*

Why InfluxDB?

Historians are Time Series Databases built by large vendors, expensive and closed source (proprietary). In contrast, Factry found InfluxDB—because it is modern, open source, offers high write throughput, has flexible retention policies, a SQL-like query syntax, and no external dependencies—to be suitable for Factry Historian and subsequently for A&S Energie needs:

Making stored time series data useful: At A&S Energie, InfluxDB is used not just to store time series data but also to make it useful. Factry encodes data about the sensor in the measurement name (they opted not to use InfluxDB's tags):

- The classical use case for a historian is trending. To make Factry Historian very easy to use for everyone at A&S Energie, Factry made it possible for the operators to simply select a metric from a drop-down menu to create graphs.
- Metrics are named after the sensor, which has a unique name from the P&ID diagram (piping and process implementation diagram). The measurement name in the database can be easily selected by plant operators.

Handling high write loads - At A&S Energie, about 2,000 measurements per second are written, and the number is even higher on the corporate level where data is aggregated from different sites.

Being platform-independent - Most historians are based on proprietary systems and are Windows-based. InfluxDB is platform-independent, and Factry runs all installations on Linux.

Providing retention policy options - Factry use InfluxDB's infinite retention policies (no load dumping) and keeps the data at its original resolution for as long as needed.

Collecting data from a local virtual machine - Since data collection in A&S Energie is done from that single plant, Factry use a local virtual machine at the site to collect the data.

Factry replaced A&S Energie's historian solution with Factry Historian. The fact that the plant's existing historian had been collecting data since 2010 required Factry to build a custom tool to migrate this data from the existing system to the new one. Despite this, data migration took roughly only a month, and once complete, the old system was switched off, and the new Factry Historian was operational.

Technical architecture

> *A good historian (and in general, any other software component) should have a simple, open API for other software to interface with.*

In replacing A&S Energie's existing process historian with Factry Historian, Factry resolved both migration and collector setup challenges that determined the project's technical architecture.

A&S Energie Data Flow

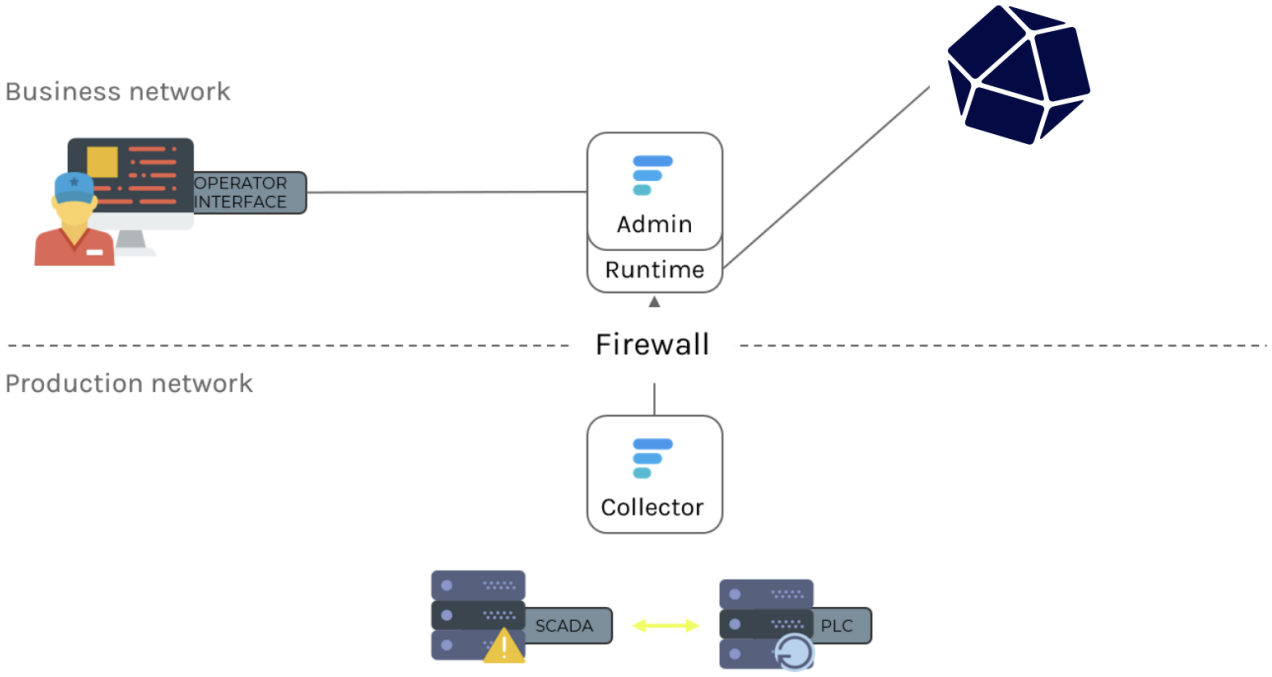

Factry's collectors collect data from PLC data systems that are controlling the process and forward it through the firewall towards an API component that sits in the business network.

This component is used to configure metrics and measurements and also has a runtime, which collects the data and forwards it to InfluxDB.

This component is used to configure metrics and measurements and also has a runtime, which collects the data and forwards it to InfluxDB.

Migrating to InfluxDB

To use Factry Historian, A&S Energie needed to migrate its data into InfluxDB. Since the plant had been running since 2010, they had years of yet unused data at second-level precision.

- Factry built a custom migration tool written in Golang to transfer the data from the old historian to the new.
- Once the tool had been built, it took Factry 47 days to migrate. The bottleneck from migration was due to the read speeds of the existing historian.
- After migration, the data was compressed from 120 GB previously (in their previous historian), to 30 GB in InfluxDB thanks to its compaction capability.
- Once Factry collected and transferred all the data, they made a switch to live data collection.
- Now, A&S Energie is collecting data from all their industrial systems through Factry Historian, and all this data is then being stored in InfluxDB.

Setting Up Collectors

In building collectors for A&S Energie, Factry needed to ensure no data is lost and needed to abide by the plant's network security policies. So Factry built collectors for different protocols, and these collectors are being used as input for InfluxDB. Collectors talk an industrial protocol on one end and HTTP on the other. They work according to a store-and-forward model (each collector has a local buffer where it stores data on-disk in case network connection goes down, and as soon as the network connection is back, the data is forwarded to the database). A local copy of the configuration is also kept, so that in the rare case that a collector starts without having any network connection, it would know which data it should collect, and can rely on its local configuration copy and start collecting data from there.

Setting up the collectors posed a challenge because the production and business networks (as is typically the case in manufacturing sites and energy plants) are separated by a firewall. Factry built collectors to ensure only outgoing connections are needed from the production network to the business network (from the machine where the collector is running to where Factry's API is collecting the data).

In between the SCADA system and InfluxDB, Factry added the following infrastructure:

- OPC-UA collector: OPC-UA is an industrial data protocol used for machine-to-machine communication in industrial settings. The OPC-UA collector is installed in the industrial part of the network. The OPC-UA collector authenticates with the backend API with a JWT token and receives its configuration from the backend API. After receiving its configuration, the collector starts retrieving data from an OPC-UA server and forwards it to the backend API, which sits in the administrative part of the network.
- A backend API, written in Golang, sits between the OPC-UA collector and InfluxDB. The backend forwards data from the collector to InfluxDB.
- A web-based frontend allows non-technical users to administer measurements (often called tags in the industry), users and collectors other than OPC-UA.

Since in this industrial context, process data is often kept for years or decades without any downsampling, Factry was able to use InfluxDB's default autogen retention policy which keeps the data forever.

Results

> *InfluxDB, combined with Grafana for data visualization, increases data visibility and puts industrial process data at people's fingertips. This increased data visibility has shown to have real value.*

Factry Historian's ease of use is enabling everyone at A&S Energie to work with the data, and A&S Energie is finally extracting value from data for trending and reporting purposes.

Trending

Trending of historical data in SCADA systems is cumbersome, and for A&S Energie in particular, data from their SCADA system was only readily available for 3 months. Anything older than that had to be placed on an external physical medium and loaded into the trending tools of the SCADA system. In contrast, with InfluxDB in place, all data is always readily available. This allowed Factry to provide Grafana to anyone in the company to create their own dashboards against data that is relevant to their specific context.

For example, the A&S operators use their custom dashboards to monitor their systems and make changes to increase plant output. The simple Grafana user interface allows them to customize their dashboards to drive even more insights from the data which leads to better plant production, better operator performance and higher work satisfaction.

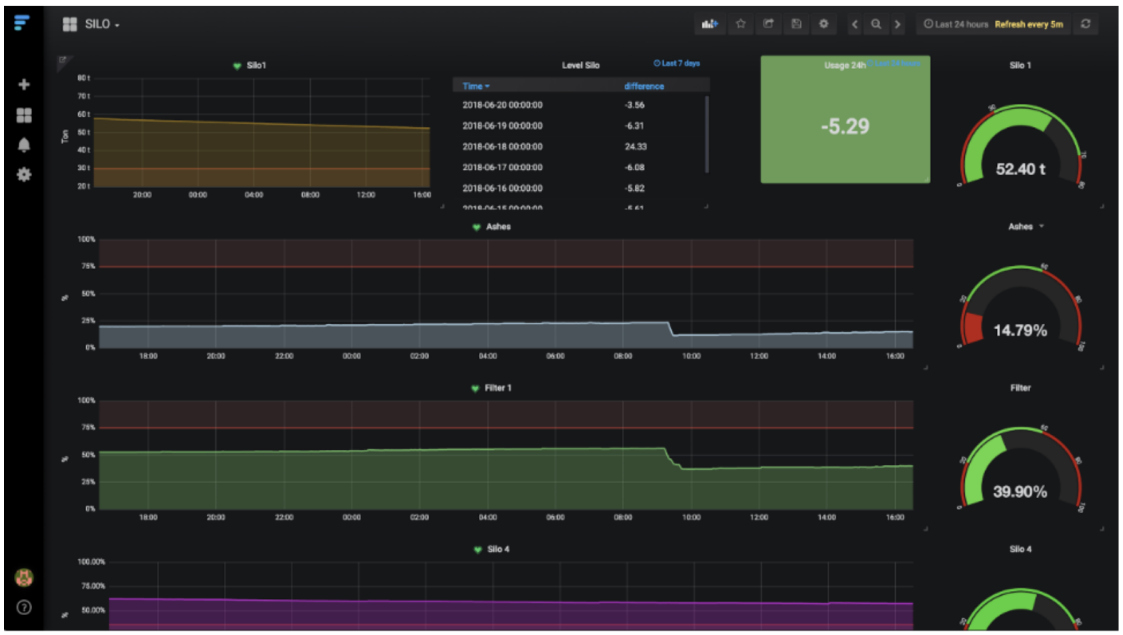

Adding to that, an early benefit that A&S gained with the new ‘transparency’ of process data was the ability to detect early signs of tank leakage, which prevented costly unplanned downtime.

Reporting

Factry is now also able to provide real-time insights into plant operations against expected performance for the management team. Data from InfluxDB is sampled at specific intervals and put into context with a PostgreSQL database which provides the necessary relations with users, equipment, etc. Management is already saving about 30 minutes per day on reporting with this new capability.

With the use of popular open source tools like InfluxDB and Grafana, Factry was able to create a solution that did not have any downtime for A&S Energie and offered seamless data migration. A&S Energie was able to continue to leverage their investment in their existing SCADA system as well as gain deeper insights into their operations, resulting in better plant production, less waste, and more productive employees.

InfluxDB and Open Source in Industry

The A&S Energie use case confirmed the viability of open source—and specifically InfluxDB—in industry, thereby confirming Factry’s founding vision of relying on open source technologies in industry and leading Factry to the following conclusions:

- InfluxDB is ready for use in industry and is really powerful in that use case.
- Open source can play a really good role in a production (and not only development) context—the project Factry completed for A&S Energie was done completely with open source components.

Factry Historian acts as a powerful, easy-to-use data collection platform for real-time data from A&S Energie’s existing production systems. With InfluxDB as Factry Historian’s storage backend, there is no limit on the amount of data that can be collected nor the amount of users or integrations.

Using Factry Historian powered by open source InfluxDB for monitoring and managing A&S Energie’s plant, Factry is fulfilling its mission of achieving data-driven operational improvement for its customers.

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.