influxdata®

AN INFLUXDATA CASE STUDY

Graphite Energy Relies on InfluxDB for Data Driven IoT Insights and Analytics

InfluxDB helps Graphite Energy maintain thousands of IoT devices in the field

FEBARUARY 2022 (REVISION 1)

Company in brief

Based in New South Wales, Australia, Graphite Energy develops Thermal Energy Storage (TES) systems that enable industrial decarbonization. Thermal Energy Storage decouples variable, intermittent, and low-cost renewable energy sources, such as wind farms or solar photovoltaic fields, from the process requirements of manufacturing plants to deliver reliable, predictable heat.

Case overview

Graphite Energy builds Thermal Energy Storage solutions designed to reduce industrial carbon footprints. The company relies heavily on data to make informed decisions about maintenance, operations, production, and research. It uses InfluxDB to collect, process, and store critical data from its devices in the field. As a result, InfluxDB plays a critical role in Graphite Energy's ability to execute its data vision.

The business challenge

Graphite Energy sought a way to reduce the carbon intensity of industrial heat production by taking advantage of low cost renewable energy. Solar and wind energy can generate variable amounts of energy from one day to the next. Graphite Energy developed its Thermal Energy Storage system to collect and store energy produced by variable, renewable sources and then reliably output that energy for an industrial process when needed in the form of heat or steam. This effectively decouples the availability of renewable energy production from the industrial requirements of that energy on the factory floor. Graphite Energy's solution allows effective, incremental displacement of fossil fuels by low-cost renewable electricity.

Technologies used:

C#, Flux, Grafana, InfluxDB, Node-Red, Python, Smartsheet

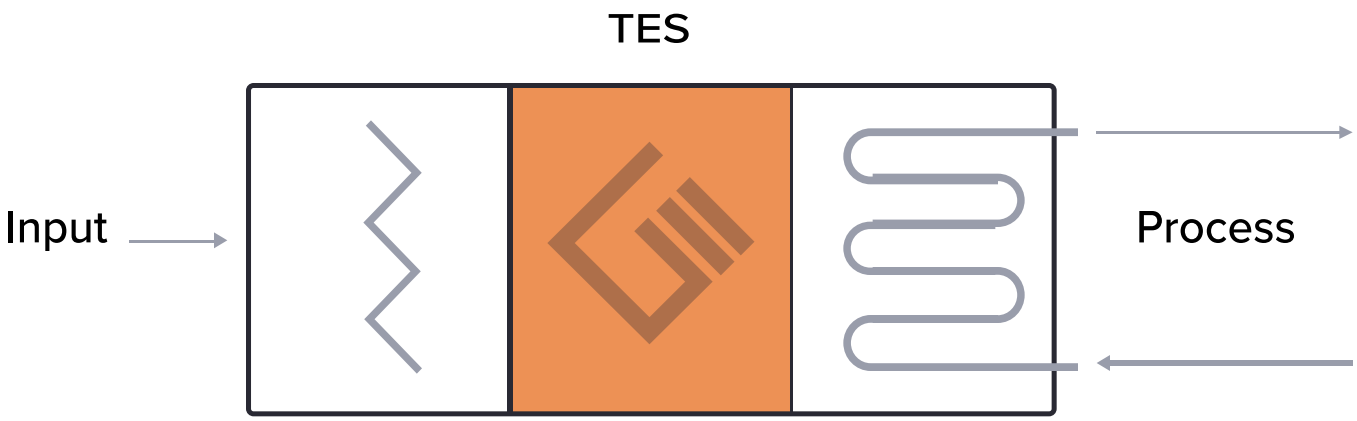

To drive efficiency and optimization of this technology, the company needs to be able to observe all its TES devices in the field and collect performance data. Taking a data driven approach to all aspects of the business enables Graphite Energy to develop products that more effectively and efficiently accelerate industrial decarbonization.

The technical challenge

Graphite Energy began its data management process in 2007 at a solar farm in Northern Germany. This installation had 2,250 independent machines, all orchestrated in real time. They used a traditional data historian to collect data. They also relied on batch-based analysis operations. For these, the optical calibration process occurred online, while all others took place offline. Each machine produced approximately ten data series, recorded every sixty seconds for a total of around 15,000 records per day.

Today, the situation is very different. Instead of monitoring a centralized machinery installation, Graphite Energy has thousands of TES devices in the field, and they log everything, such as the internal temperature of the storage medium and power consumption. Data collection occurs both locally on each device and in the cloud, and gets stored in InfluxDB Cloud, a time series database. They rely on a range of cloud-based and Edge-compute analytics that they visualize in real time using web dashboards. Each machine produces approximately 100 data series that get recorded anywhere from every 1–15 seconds for a total of around 1,000,000 records per day.

With so many devices in the field, collecting all this data, the Graphite Energy team needed the ability to install software on each machine to manage data collection. They wanted to control these devices remotely and sought something that was both reliable and easy to use. Having used Microsoft SQL and Azure Time Series Insights in their production environment (and even more time series databases in testing), they settled on InfluxDB as their time series database.

The solution

Graphite Energy's time series data architecture has several components. On the customer end, they have a local control network, a machine data interface, an operator interface, and some edge compute on each device. There's also a local instance of InfluxDB on each unit. They built a remote operations center to support their customers, and which houses various dashboards and engineering tools (as well as another InfluxDB instance). All this connects to their primary data store an instance of InfluxDB Cloud, where additional processing occurs.

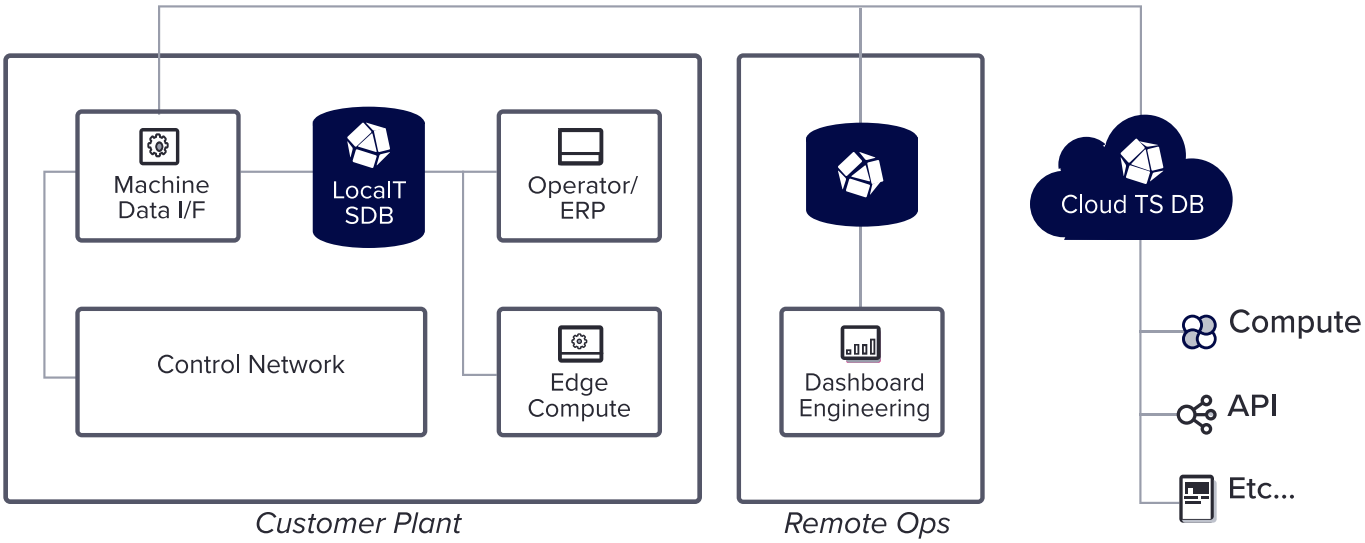

Graphite Energy's current data management solution centers on a Node-RED workflow. They use an industrial controller made by Weidmuller that includes a real-time logic controller system with a Node-RED instance on the side. Graphite Energy uses this on-board Node-RED instance for data collection. At a basic level, Graphite Energy uses Smartsheet to grab configuration information, collects the data from the TES devices using Node-RED, engineers and processes that data with a variety of tools (e.g., Python, C#), and then stores it in InfluxDB.

The different nodes of the Node-RED workflow allow Graphite Energy's engineers to amend and process the data as it works its way through to transform it into a consistent shape for storage and analysis in InfluxDB. Graphite Energy's main development and operation workflow relies on InfluxDB and Grafana. They use Grafana extensively for internal and customer dashboards, and to do some simple automation tooling.

Once the data is in InfluxDB, Graphite Energy's engineers use Flux to query and transform that data to feed their digital twin feature. Flux allows them to query, combine, and transform data to make it more useful.

> *"We really, really like Flux...We have started to do some nice things at query time that traditionally we would have done later in the toolchain...Flux makes it very easy. We haven't been able to do this in other query languages."*

Byron Ross, COO, Graphite Energy

Results

The digital twin is the culmination and end-product of Graphite Energy's time series data management process. They use time series data to create a real-time digital model of a TES unit that is accurate to within about 5% of actual machine performance. Digital twin lets Graphite Energy roll forward and backwards in time to track device performance and is becoming a very powerful part of their predictive toolkit for production optimization.

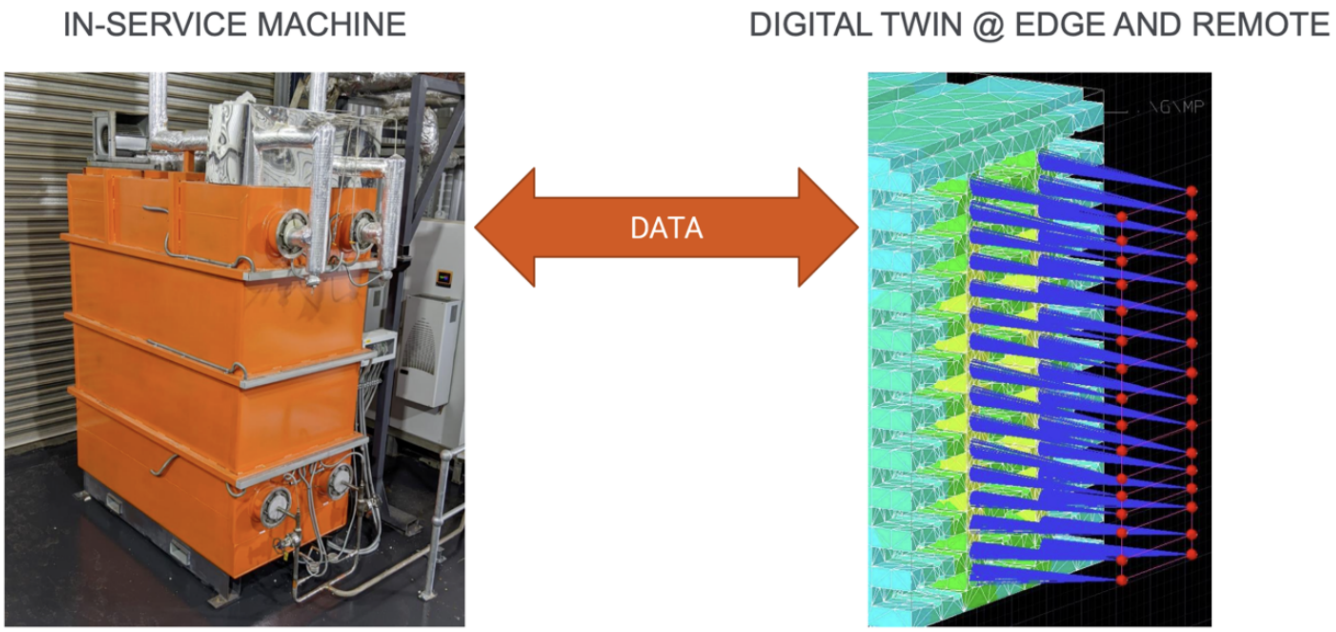

> *We are absolutely sticking with InfluxDB for our data storage. This is enabling our success and, we think, enabling our customers' success."*

Byron Ross, COO, Graphite Energy

What's next

Graphite Energy continues to expand its use of time series data. The company anticipates scaling its solution to accommodate 100x more series than they currently collect. Much of this is due to an increased number of sensors on their machines in the field and a higher number of actuators. There are several peripheral systems that they can collect data from but have not yet included in their process. Folding this data into the process will account for a portion of the anticipated data volume increase. They are also moving towards putting some on-device machine learning on the industrial controller using Python to further support their digital twin and edge compute strategies. COO Byron Ross noted that, "The on-premises and InfluxDB Cloud instances are really going to help us achieve our goals there."

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.