

FROM TELEMETRY DATA POINT TO THE ENERGY TWIN: AN OT-EDGE METHODOLOGY AND SEMANTIC CONTEXTUALISATION FOR COST REDUCTION IN THE CERAMICS INDUSTRY

Valter Cartella

Freelance consultant in IT/OT Security and Industrial Digitalisation — Spain

ABSTRACT

The ceramics industry is characterised by energy-intensive processes, in which kilns, dryers, compressors, production lines and auxiliary systems have a significant impact on operating costs, competitiveness and decarbonisation targets. However, in many industrial settings, energy telemetry is still treated as a set of consumption graphs that are detached from the factory's actual operational context. The result is often a 'data graveyard': large volumes of measurements are available, but they are difficult to translate into decisions, priorities for action and energy-saving measures.

This paper proposes a lightweight OT-Edge methodology to move from traditional energy monitoring towards an Energy Twin applicable to the ceramics industry. The methodology is based on the acquisition of energy and operational data at edge level, data transmission via an MQTT broker, semantic fusion and contextualisation using edge processing logic, and data storage in a time-series database for the generation of KPIs, analysis and operational insights.

The distinctive feature of the proposal is the semantic contextualisation of an energy data point. Consumption is not analysed as an isolated value, but in relation to the asset, line, area, machine state, data point quality and, where necessary, a lightweight input from the operator via a tablet or simplified HMI interface. This allows energy consumption to be correlated with states such as normal production, set-up, format changeover, breakdown, maintenance or standby, making it possible to calculate unproductive consumption, energy losses, anomalies and optimisation priorities.

The methodology is designed as a 6–8-week, non-invasive, low-risk OT-Edge proof of concept (PoC), aimed at rapidly validating whether contextualised energy data can generate operational value before scaling up to a more comprehensive Energy Twin. The proposal is supported by an end-to-end technical demo that has already been validated in a laboratory environment, based on a simulated industrial asset, JSON telemetry, MQTT, Node-RED, InfluxDB and Grafana.

The main contribution lies in proposing a pragmatic approach from a telemetry data point to the Energy Twin: not a complete and immediate Digital Twin, but a contextualised industrial database, useful for reducing costs, identifying opportunities for cost savings, and preparing for future advanced digitalisation initiatives.

1. INTRODUCTION

In the ceramics sector, energy is not merely an operating cost: it is a strategic factor that influences competitiveness, quality, efficiency and decarbonisation. Processes such as firing, drying, pressing, glazing, internal handling, air compression and packaging require significant amounts of energy and contribute directly to the company's operating costs.

In this context, industrial digitalisation cannot be viewed merely as a technological trend. It must become a practical tool for better understanding the production process, reducing losses, improving the quality of decision-making and enhancing competitiveness. Data, if collected, contextualised and used correctly, can become a real competitive advantage.

However, in practical industrial terms, there is often a gap between data availability and effective use. Many companies have access to telemetry, energy measurements, production history, machine events and alarms, but such data often remain isolated, unstandardised or difficult to correlate. As a result, energy monitoring risks being limited to the display of consumption curves, without automatically generating operational insights.

This paper proposes an **OT-Edge methodology** to bridge that gap. The aim is to transform raw energy telemetry into contextualised industrial data, linking energy consumption to the actual operational context: asset, line, area, machine state, process phase, data point quality and, where necessary, operator input.

The intended outcome is not the immediate creation of a full Digital Twin, but rather the development of an initial-level Energy Twin: an energy-operational representation of the facility capable of supporting analysis, KPIs, the identification of losses and prioritisation of actions.

2. THE PROBLEM: THE ENERGY DATA GRAVEYARD

In many factories, the problem is no longer simply a lack of data. Often, the problem lies in the difficulty of transforming available data into useful decisions. Kilns, dryers, compressors, production lines and auxiliary systems can generate large volumes of telemetry, but this telemetry is not always linked to the operational context that determines its meaning.

An energy consumption figure, taken in isolation, provides limited information. Knowing that an asset is consuming 120 kW may be useful, but it is not enough to understand whether that consumption is productive, unproductive, abnormal or avoidable. The same figure takes on very different meanings depending on whether it relates to normal production, set-up, format changeover, waiting for materials, maintenance, breakdown or standby.

Traditional energy monitoring allows for consumption to be observed but does not always enable interpretation. A dashboard can show trends in power or energy consumption but without operational context, it is difficult to answer key questions:

- How much energy is consumed during normal production?

- How much energy is consumed during set-up and format changeovers?
- Which assets consume energy even when they are not producing?
- Which stoppages result in unproductive energy consumption?
- Which energy anomalies recur over time?
- Which energy-efficiency measures should be given financial priority?

This phenomenon can be described as a 'data graveyard': data that are available but are not being fully leveraged. The proposal put forward stems precisely from this critical issue: measuring energy does not yet mean understanding it.

3. FROM ENERGY MONITORING TO THE ENERGY TWIN

The transition from energy monitoring to the Energy Twin requires a conceptual leap. It is not simply a matter of collecting more data or visualising the data using more sophisticated dashboards. It is about building a coherent picture of the facility's energy and operational behaviour.

From this perspective, energy monitoring represents the first level: observing how much energy is consumed. The next level is operational energy intelligence: understanding the context in which such consumption occurs. The Energy Twin represents a further step: a database and analytics platform capable of describing the energy behaviour of assets, lines and processes over time, linking consumption, operational state, production conditions and opportunities for improvement.

The Energy Twin should not be confused with a fully-fledged industrial Digital Twin. A mature Digital Twin may include physical models, complex simulations, predictive maintenance, quality control, production, logistics and multivariable optimisation. The Energy Twin, by contrast, can be understood as a pragmatic and focused first step: an energy-operational twin designed to understand how, when, where and why energy is consumed.

This approach is particularly well suited to the ceramics industry, where energy consumption is high and where a better understanding of the relationship between energy, the production process and the operational context can yield tangible benefits.

4. PROPOSED OT-EDGE METHODOLOGY

The proposed methodology is based on a lightweight OT-Edge architecture, designed to be initially implemented on a single asset or pilot line, without requiring a complete factory-wide transformation.

The principle is simple: collecting data close to the process, contextualising the information at the edge, storing it in a consistent manner, and transforming it into KPIs that are useful for decision-making.

The main stages of the methodology are as follows:

1. Selecting an asset or line with a high energy impact;

2. Identifying available data from PLCs, meters, SCADA, MES or other systems;
3. Acquiring energy and operational telemetry at edge level;
4. Publishing the data via an MQTT broker;
5. Acquiring, where necessary, semantic input from the operator via a tablet or lightweight HMI;
6. Standardising and fusing the data at edge level;
7. Storing the enriched data in a time-series database;
8. Calculating energy and operational KPIs;
9. Identifying losses, anomalies and unproductive consumption;
10. Defining a roadmap towards an Energy Twin and an industrial Digital Twin.

The proposed architecture is designed to be passive and non-intrusive. The aim is not to interfere with process control, but rather to observe, collect data, correlate and analyse. This aspect is particularly important in industrial environments where safety, production continuity and operational reliability are priorities.

5. SEMANTIC CONTEXTUALISATION: THE DISTINGUISHING FACTOR

The central element of the methodology is semantic contextualisation. An energy data point is not regarded as an isolated value, but as part of a broader industrial event.

A contextualised energy data point may include:

- Timestamp;
- Plant;
- Area;
- Line;
- Asset;
- Asset type;
- Machine state;
- Operator context;
- Power;
- Cumulative energy;
- Current;
- Voltage;
- Data quality;
- Data source.

There is a significant difference between raw data and contextualised data. A value such as '120 kW' merely indicates a measurement. Information such as 'kiln_01, line_01, normal production, 120 kW, good data point quality, accurate timestamp' makes it possible to attribute that consumption to a specific operating state.

This logic makes it possible to answer far more useful questions:

- What is the energy cost of a format changeover?
- How much energy is consumed during standby modes?
- Which assets are consuming energy at levels that are unusual for their operating state?
- Which operational phases are causing recurring losses?
- Which measures can generate savings most quickly?

In this context, semantics is not an abstract concept. It is the link between a technical data point and operational reality.

6. SEMANTIC INPUT FROM THE OPERATOR

A distinctive feature of the proposal is the ability to incorporate minimal operator input. In many industrial facilities, not all operational states are available automatically, reliably or with sufficient granularity. Some of the information remains within the operators' practical knowledge or is recorded in an unstructured manner.

For this reason, the methodology utilises a tablet, smartphone or other lightweight HMI interface as the operator's semantic input layer. The interface may feature a small number of clear operational states, for example:

- Normal production;
- Set-up / format changeover;
- Breakdown;
- Maintenance;
- Idle / standby;
- Other, with a text note.

When the operator selects a state, a timestamped event is generated which can be published via MQTT. This event is subsequently fused with the energy telemetry at edge level.

The aim is not to replace PLCs, SCADA or MES systems. On the contrary, operator input serves as a supplementary mechanism, which is particularly useful in the early stages of a proof of concept (PoC), in legacy installations or in contexts where operational data are not yet fully available.

This mechanism makes it possible to quickly transform a raw energy measurement into an enriched data point. For example, the system does not merely record that a compressor consumes 12 kW but also indicates whether those 12 kW are associated with a set-up, standby, maintenance or normal production state.

7. REFERENCE ARCHITECTURE OF THE DEMO STATION

The methodology has been validated using an OT-Edge Demo Station in a laboratory environment. The Demo Station does not represent a final product or a solution linked to a specific supplier, but rather a technical laboratory designed to demonstrate the feasibility of the method.

The validated chain consists of:

- A simulated industrial asset using the Python Publisher;
- JSON telemetry;
- The Mosquitto MQTT broker;
- Node-RED for receiving, transforming and fusing data;
- InfluxDB for time-series storage;
- Grafana for visualising KPIs and dashboards.

The simulated asset generates telemetry relating to an industrial compressor, including fields such as power, energy, current, voltage, machine state, asset, line, area, data quality and source.

The architectural diagram of the methodology shows two distinct data flows that converge on the same MQTT broker:

1. The flow of energy and operational telemetry data from assets, PLCs, meters or simulators;
2. The flow of operational context data from tablets or operator interfaces.

Downstream of the MQTT broker, Node-RED acts as a fusion layer, providing semantic contextualisation. The enriched data are then stored in InfluxDB and used to calculate KPIs and operational energy insights.




Figure 1. Reference architecture of the proposed methodology. Two distinct data streams – energy and operational telemetry from industrial assets, and operational context entered by the operator via a lightweight interface – converge at the MQTT broker. Downstream, Node-RED performs the semantic fusion and contextualisation of the data, which are subsequently stored in InfluxDB to generate KPIs, energy cost analyses and a database geared towards the Energy Twin.

8. KPIs AND COST REDUCTION OPPORTUNITIES

The value of the methodology becomes apparent during the analysis phase. Once the energy data point has been put into context, it is possible to calculate KPIs that are more useful than simply visualising consumption.

Examples of energy-related operational KPIs include:

- Energy consumption by operating state;
- Energy consumed during normal production;
- Energy consumed during set-up and format changeovers;
- Energy consumed in standby or waiting mode;
- Idle energy loss;
- Energy waste index;
- kWh per unit produced;

- Abnormal consumption per asset;
- Comparison between lines or shifts;
- Energy cost per operating stage;
- Economic priority for action.

The key point is that savings stem from the ability to attribute consumption to the correct operating condition. Knowing that a line consumes a certain amount of energy is useful. Knowing how much of that energy has been consumed in production, set-up, idle time or downtime is much more useful.

This capability makes it possible to distinguish between productive and unproductive consumption, identifying opportunities for optimisation that would otherwise remain hidden.

In contexts such as Spain, the availability of reliable, historical and contextualised energy data can also facilitate the technical assessment of efficiency measures and their eventual integration with regulated mechanisms for recognising energy savings, where applicable. Whilst this is not the primary objective of the methodology, it may represent an additional benefit when the collected data enable improvement measures to be documented more robustly.

9. POC PROCEDURE: 6–8 WEEKS TO VALIDATE ITS BENEFIT

The proposed methodology can be implemented through a lightweight and transferable proof of concept (PoC), designed to reduce complexity, risk and validation times.

The aim of the PoC is not to digitise the entire factory, but to demonstrate quickly whether contextualised energy data can lead to operational decisions and real opportunities for savings.

A possible implementation procedure comprises the following stages:

1. Selection of an asset or pilot line;
2. Preliminary analysis of available data;
3. Definition of a minimum taxonomy of operational states;
4. Temporary installation of the OT-Edge architecture;
5. Acquisition of energy and operational telemetry;
6. Acquisition of the automatic or manual operational context;
7. Semantic fusion at edge level;
8. Time-series storage;
9. Calculation of KPIs;
10. Analysis of savings opportunities;
11. Final report and scaling roadmap.

The **indicative duration of the PoC is 6–8 weeks**. That is long enough to collect representative data, observe different operating states, validate telemetry quality, calculate preliminary KPIs and identify potential opportunities for optimisation.

The following table summarises a possible step-by-step roadmap:

Phase	Indicated duration	Aim	Risk
PoC OT–Edge on asset or pilot line	6–8 weeks	Validating data acquisition, semantics and KPIs	Low
Scaling to line or process	2–3 months	Correlating energy, operating states and production	Medium–low
Operational Energy Twin	3–6 months	Modelling energy-operations behaviour	Medium
Industrial Digital Twin	Later stage	Integrating processing, quality, maintenance and simulation	Higher

This roadmap helps to avoid a common pitfall in industrial digitalisation projects: starting with initiatives that are too broad, costly or difficult to validate. In contrast, the methodology proposes an incremental approach: starting small, validating the value, measuring the benefits, and only then scaling up.

10. APPLICABILITY TO THE CERAMICS INDUSTRY

The methodology is particularly well suited to the ceramics sector because it can be applied progressively to different assets and processes. This is particularly relevant for the Castellón ceramics district, where industrial competitiveness increasingly depends on optimising energy use in kilns, dryers, production lines and auxiliary systems.

Examples of applications include:

- Kilns;
- Dryers;
- Compressors;
- Pressing lines;
- Glazing lines;
- Decorating lines;
- Internal handling and internal logistics;
- Packaging;
- Auxiliary systems.

In each of these contexts, energy consumption can be analysed in relation to operating states, process stages, formats, batches, quality, maintenance and production conditions.

A specific example is a format changeover. During a format changeover, a production line may continue to consume energy even though it is not directly producing value. Without

contextualisation, this consumption may appear to be a normal part of the energy curve. With semantic contextualisation, however, it is possible to quantify the energy cost of format changes, compare it across shifts or production lines, and identify opportunities for reduction.

Another example relates to compressors and auxiliary systems. In many facilities, these assets continue to consume energy even when operating at low capacity, in standby mode or inefficiently. Linking energy consumption to operating state makes it possible to identify recurring losses and set priorities for action.

11. SECURITY, NON-INTRUSIVENESS AND DATA GOVERNANCE

In an industrial setting, any digitalisation proposal must take into account security, operational continuity and data governance. For this reason, the proposed methodology adopts a non-intrusive approach, focusing on observation and analysis, without interfering with process control.

The OT-Edge PoC is designed as a local, controlled architecture, connected to industrial assets in read-only mode where technically feasible, and designed to collect energy and operational telemetry without sending commands to PLCs, controllers or automation systems. The aim is to monitor the facility's energy and operational behaviour, not to alter its operation.

From a segmentation perspective, the solution can be implemented as a small, dedicated internal network that is isolated, or virtually isolated, from the existing control network. In this configuration, the MQTT broker, the edge processing logic, the time-series database and the dashboards remain within a controlled local perimeter, accessible only to authorised personnel, such as the consultant, the operator or the technical manager in charge of the PoC.

This model reduces the surface area exposed to risk and ensures that the Demo Station is compatible with a prudent approach to industrial security. The solution does not require a mandatory cloud connection, does not introduce remote control functionality over industrial assets, and can be managed as a temporary analysis environment, separate from critical automation systems.

In terms of the Purdue Model, the architecture must respect the industrial zone in which it is located, without creating ungoverned paths to higher or lower levels. Integration with PLCs, SCADA, MES, cloud platforms or enterprise systems can only be assessed at later stages, using segmentation rules, access control, identity management, logging and data governance consistent with the company's industrial cybersecurity policies.

The methodology does not impose a specific final architecture, but proposes an initial, replicable, and secure-by-design model suitable for rapid validation. In later phases, the same logic can be integrated with SCADA, MES, quality management systems, maintenance systems, cloud platforms, data lakes or AI tools, depending on the organisation's digital maturity and security requirements.

12. CONCLUSIONS

The proposal presented here introduces an OT-Edge methodology for transforming energy telemetry into contextualised operational insights. The key focus is not on measuring more data, but on attributing meaning to data that are already available or easily obtainable.

In the ceramics industry, where energy, quality, productivity and competitiveness are closely linked, this contextualisation can generate tangible value. Linking energy consumption to assets, production lines, operating states, process stages and operator input makes it possible to identify unproductive consumption, anomalies, losses and priorities for action.

The proposed methodology does not aim to create a full-scale Digital Twin straightaway. Instead, it proposes a concrete, measurable and low-risk first step: a 6–8-week OT–Edge proof of concept (PoC) to validate whether the semantic contextualisation of the energy data can provide support for operational decisions and identify energy-savings opportunities.

The Energy Twin thus serves as a practical bridge between traditional energy monitoring and the industrial Digital Twin. A well-structured energy and operational database can power not only dashboards and KPIs, but also future models for anomaly detection, process optimisation, maintenance, quality, decarbonisation and strategic planning.

In short, the proposal aims to demonstrate that the Energy Twin should not be seen as just another technological trend, but rather as a pragmatic methodology for transforming contextualised energy data into a real, measurable and sustainable competitive advantage for the ceramics industry.

13. REFERENCES

- [1] IEC 62264 / ISA-95, Enterprise-control system integration.
- [2] IEC 62443, Industrial communication networks — Network and system security.
- [3] ISO 50001, Energy management systems — Requirements with guidance for use.
- [4] MQTT, OASIS Standard, Message Queuing Telemetry Transport.
- [5] Industrial Internet Consortium, Industrial Internet Reference Architecture.
- [6] Documentación técnica de referencia sobre bases de datos time-series, edge computing y contextualización de datos industriales.