Closing The Loop

A Field Report: Agentic AI in Legacy Buildings

Quantum Spark
CONSULTING

AT A GLANCE

THE BRIEF

We were tasked to save energy spent in a 1960s office block and prove that we could make a return on investment in under three years. The building had an ageing heating system, single-glazed and draughty windows, and many complaints from staff that it was either freezing cold or roasting. We decided to tackle the thermostatic radiator valves and replaced them with smart versions. Everything was connected through a couple of wireless base stations across the building. This enabled an AI Agent, whose goal was no longer to keep every room warm all day, but to maintain comfort precisely when, and where, it was needed. With room-by-room compensation for weather and occupancy, some of the results were quite surprising...

ANALYSIS AND IMPLEMENTATION

+ Rather than rushing to install new controls, the first step was to understand the building as it actually was.
+ Once we had our baseline, we switched the old manual valves for intelligent wireless versions, but these were deliberately left in manual mode; every adjustment an occupant made became a data point describing how people were interacting with the building.
+ After a period of observation, we enabled the building's Agent (AI) to begin making decisions for itself. Initial heating schedules were introduced, tailored not to a single timetable for the whole building but to how each space was actually used.

RESULTS

25% total energy saving, across both electricity and gas.

100% investment returned within year one.

74% fewer manual adjustments, people were happier.

"The smartest move wasn't teaching people about the heating... it was teaching the heating about people."

Richard Lansdowne
Director Quantum Spark Ltd

CLOSING THE LOOP

Facilities departments often already have more data than they know what to do with. Meters, sensors, dashboards: the last decade of smart building technology delivered plenty of numbers and graphs, but fell short, with comparatively few decisions. Call this the open loop: information flows in, but something outside the system (a person reacting to an alert or a simple rule that was fixed months ago) has to close the loop and take the action. This gap, between analysis and decision, is where IoT has failed to deliver much of the promised return on investment.

What follows is an account of a project where we used an Agent to 'close the loop'. This is not about the smart hardware; the sensors and wireless valves involved are the same devices that have been available for years. What's different is the "system" that reads what's happening, decides what to do about it, acts, and then checks its own work, continuously, without anyone rewriting the rules.

Example 1960s office building

The building and ROI target

When Quantum Spark was asked to investigate an office building with a reputation for being impossible to heat properly, the question was asked "what can AI do for us here?" And, of course, how much cost can we save?

The brief was straightforward, if familiar to anyone managing an ageing estate. The client knew they had a problem: staff complained that offices were either freezing or unbearably hot, often within the same working day, energy bills were climbing, and confidence in the heating system had long since evaporated. New radiators had been fitted, thermostatic valves were in place, yet the problems persisted. The target set was equally clear: demonstrate a return on investment within three years.

It would, in the end, do rather better than that.

Our first visit took place on a cold winter afternoon, the temperature outside hovering between four and six degrees. Inside, the picture was oddly inconsistent: unused spaces had been left heated needlessly, while people in occupied offices described spending every morning trying to get warm before their own office became unbearably hot later in the day. That contradiction was the first clue: this wasn't a boiler issue; something more interesting was going on.

Walking through the offices confirmed it. Every radiator was fitted with a traditional thermostatic valve, the kind with a dial that runs from one to five where no one knows what those numbers actually mean, apart from a rough sense of 'hotter' or 'colder'. People had turned the radiators into a set of individual experiments: some fully open, some almost closed, and neighbouring desks in the same office running at entirely different settings. The building was, essentially, the sum of hundreds of small, uncoordinated attempts to compensate for a system that never behaved the way anyone expected or desired.

Establishing a baseline

Rather than rushing to install new controls, the first step was to understand the building as it actually was. Energy meters went in to establish a genuine baseline of usage, while weather data (external temperature, wind speed and direction, rainfall, solar radiation) was included to enable the system to learn the building's response to the physical conditions.

What emerged was telling: while the weather clearly influenced internal temperatures, gas consumption stayed remarkably constant. Regardless of the weather-based demand, this building appeared to be merely heated by the clock. The sun was warming south-facing offices on sunny afternoons and ill-fitting windows were stripping heat from exposed rooms on windy days, while the heating system itself appeared oblivious.

Original inefficient valves

New smart valves

Replacing the valves and valuable extra data

Once we had our baseline, we switched the old manual valves for intelligent wireless versions, but these were deliberately left in manual mode. Occupants kept using them exactly as before; nobody asked staff to change a thing. The only real change was that new valves displayed temperature in degrees rather than an arbitrary number, and every adjustment an occupant made became a data point describing how people were interacting with the building.

The change to temperature-based valves produced an immediate, moderate drop in energy use but, more usefully, exposed patterns that were previously invisible. Radiators in the same office were often set at dramatically different temperatures, some occupants adjusting the heating repeatedly throughout the day while others never touched it. A slight complication was discovered: the valves' internal thermostat couldn't be trusted to measure room temperature accurately because the heating pipes ran beneath the radiators. Even a switched-off radiator sat above pipework carrying hot water passing through to other rooms, distorting its own reading. We deployed independent room sensors to give the system a proper reference point: not what the valve believed the temperature to be, but what people were actually experiencing. Thus, the system could learn the underlying effect of the heating pipes and compensate for those changes too.

Closing the loop; with surprising results

After a period of observation, we enabled the building's Agent (AI) to begin making decisions for itself. Initial heating schedules were introduced, tailored not to a single timetable for the whole building but to how each space was actually used: offices warmed ahead of the working day with start times shifting by the weather's impact on each specific room's thermal behaviour. Meeting rooms were on a different pattern again while training rooms were aligned to bookings.

The Agent's aim was no longer to keep every room warm all day, but to maintain comfort when, and where, it was needed. Hence, gas consumption fell sharply almost at once… however, that was only part of the story. Something else changed that had nothing directly to do with gas: electricity consumption fell by almost exactly the same proportion as gas.

For years, occupants appeared to have run a routine nobody had ever questioned: arrive cold, reach for the radiator dial, switch on a portable heater under the desk, then reverse the whole process once the office became too warm. With the Agent ensuring that the building reached a comfortable temperature before people arrived, that routine simply stopped, not because anyone told staff to turn the heaters off, but because nobody felt the need to switch them on. The system had no connection whatsoever to those heaters; it could neither see nor control them, but solving the comfort problem on arrival removed the need to use them.

> "The system didn't take people out of the loop. It took away the reason they kept being pulled back into it."

A third pattern showed up in the data: manual adjustments to the radiator valves fell by seventy-four per cent, not through any policy telling people to leave the controls alone, but because they stopped feeling the need to intervene.

By this stage the original business case was already comfortably exceeded: energy consumption down by around twenty-five per cent, with the capital cost recovered in under a year. The most interesting phase of the project, though, was only just beginning.

What makes it agentic

Unlike a conventional building management system, this one was never designed to stop learning once installation finished. It kept watching and 'listening' to its users. Every room's response was compared to its own predictions, and constantly refined, improving its understanding of the building: which offices overheated in the low winter sun, which stayed stubbornly cold from draughts around ageing window frames, which radiators seemed underperforming or partially blocked, how quickly different parts of the building recovered after a cold night, and how that recovery time shifted through the seasons.

None of those things were programmed by an engineer. They emerged because the Agent in the system kept asking itself the same question, over and over: what should I do differently next time to achieve my goal more effectively? That question, repeated continuously, on a room-by-room, office-by-office basis, and acted on without waiting for a person to rewrite the rules, is the practical difference between automation and what's now called Agentic AI.

> "Automation follows a rulebook someone else wrote whereas an AI agent is given a goal to achieve. It works out for itself, continuously, how best to reach the goal as the seasons and the people using it keep changing."

It's worth being precise about what actually changed here, because it's less than the word AI tends to imply. The building is still heated by ordinary radiators. The sensors are little different from those fitted in countless commercial buildings, and the valves look much like any others. What's different is that the building no longer waits to be told how to behave. It observes, decides and acts continuously, in something like the manner of an experienced facilities manager, only checking in every few minutes rather than once a week, and never getting distracted.

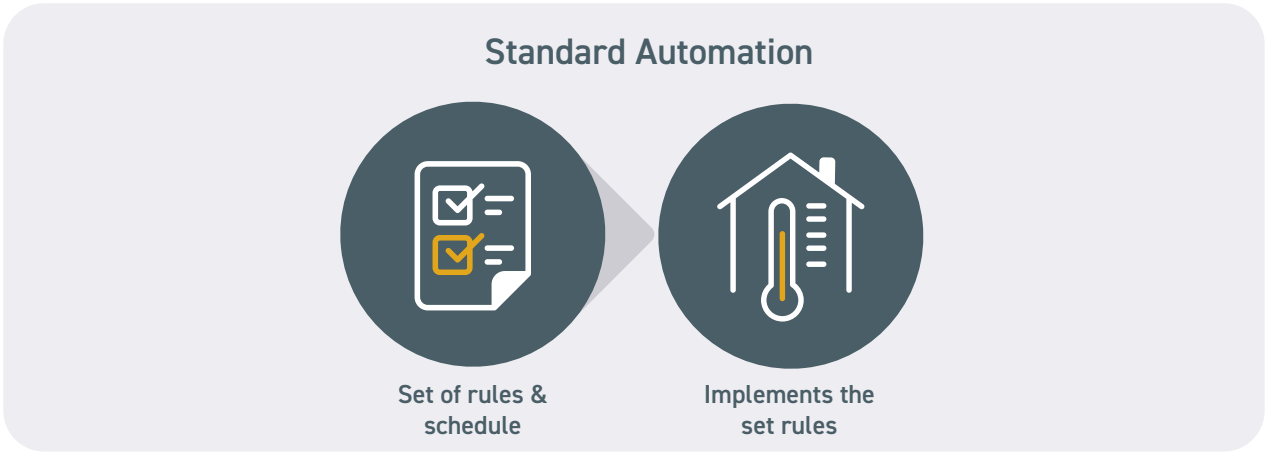

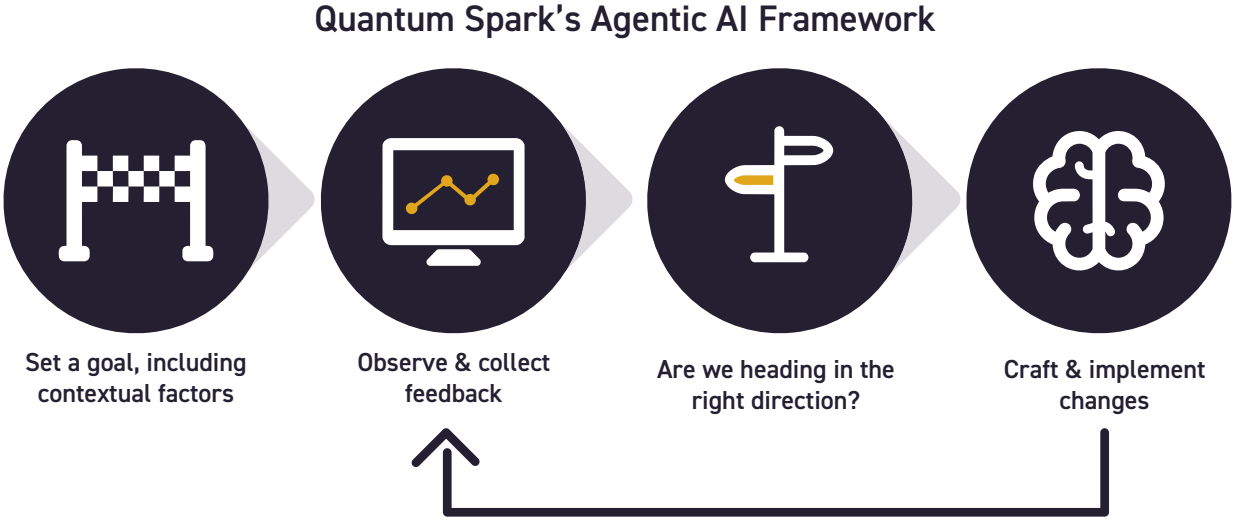

Key takeaways

For any organisation weighing up a similar investment, three things from this project are worth holding onto.

1. **Observation has to come before optimisation.** It would have been easy for us to skip straight to intelligent scheduling and miss the time spent simply watching how the building and its occupants behave, but without that baseline there would have been no way to tell how much of the eventual saving came from the technology and how much from a mild winter.

2. **The most valuable savings are not always the ones anyone went looking for.** Nobody set out to reduce electricity consumption in this project; the Agent simply removed the reason why people switched on their portable heaters in the first place. The lesson generalises well beyond heating: get the primary problem right, and secondary costs often take care of themselves.

3. **What made the difference here was a system with a standing goal rather than a fixed rulebook, continuously comparing its own predictions against what actually happened and correcting course.** This is the distinction that gives this piece its title. A fixed schedule, however cleverly configured at the outset, will always drift out of step with a building's actual life as the season changes, the occupancy pattern shifts, or a radiator gradually starts to clog, and someone eventually has to notice and step in by hand. That is the open loop, and it is where most building technology still sits. What made all the difference is the loop closing itself, and also, in the end, is where the return on investment actually came from.

It's worth noting that most of what this system does has a familiar name in any other context. Watching, learning, adjusting to the people around you: that's 'reading the room'; merely applied to a building rather than a person.

Closing The Loop
A Field Report on
Agentic AI in Buildings

Published
August 2026

Quantum Spark Ltd,
Winchester, UK

quantumsparkltd.com
07920061190
info@quantumsparkltd.com