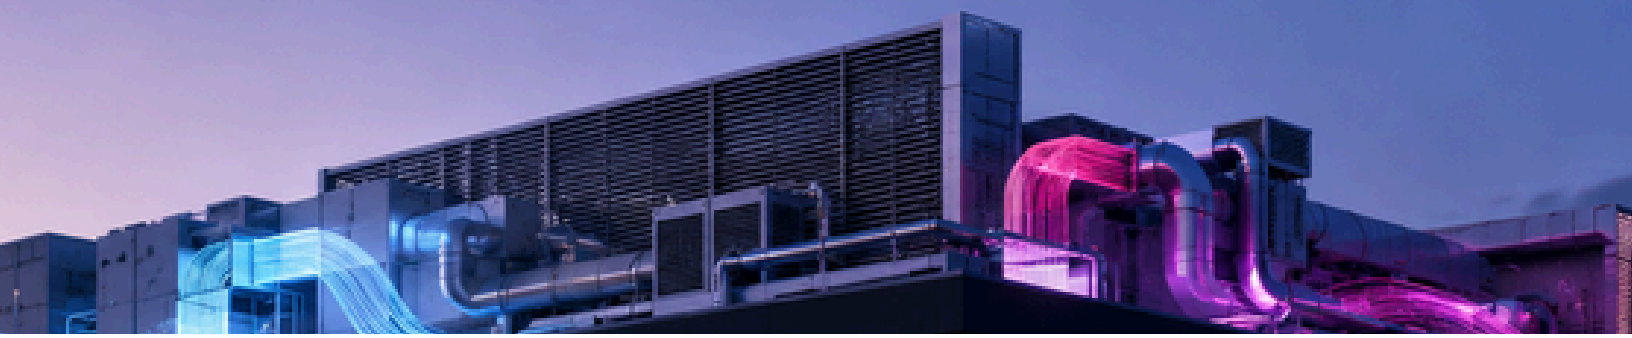

SOLUTION BRIEF

20% Lower Energy Costs Without Sacrificing Comfort

How Adaptive AI Is Changing HVAC Control

Executive Summary

Energy consumption in commercial buildings represents one of the largest controllable operating expenses for facility owners and portfolio managers. HVAC systems alone account for roughly 40% of a building's total energy use, yet most are still governed by static, rule-based controls that cannot adapt to real-time conditions. The result is a persistent trade-off between energy efficiency and occupant comfort, one that traditional building management systems have never adequately resolved.

Eng IndX, in partnership with AgileRL, has developed and deployed an adaptive HVAC control system powered by Reinforcement Learning (RL). This solution replaces reactive threshold logic with an intelligent agent that learns and optimizes in real time, balancing energy reduction against thermal comfort requirements. In validated deployments, the system delivered 20% greater energy savings, 13% improvement in comfort stability and 35% faster deployment timelines compared to conventional approaches.

This brief outlines the business case, the technical approach and the measurable results for building operators, energy managers and real estate decision-makers evaluating intelligent building controls.

The Business Problem

Large commercial buildings face a structural problem in how they manage heating and cooling. Outdoor temperatures fluctuate throughout the day. Solar gain shifts as the sun moves. Occupancy patterns are irregular and often unpredictable. The building envelope itself introduces thermal inertia, meaning that control decisions made now produce effects that only become visible minutes or hours later.

Conventional HVAC controls rely on fixed setpoints and threshold-based rules. These systems react to conditions after they have already changed, rather than anticipating what comes next. In practice, this means facility teams are forced to choose between two suboptimal outcomes: **overconditioning** spaces to avoid complaints, which drives up energy costs, or **underconditioning** to save energy, which erodes occupant satisfaction and, in tenant-occupied buildings, can directly impact lease retention.

Neither outcome is acceptable for organizations under pressure to reduce operational costs, meet sustainability commitments and maintain high-quality indoor environments.

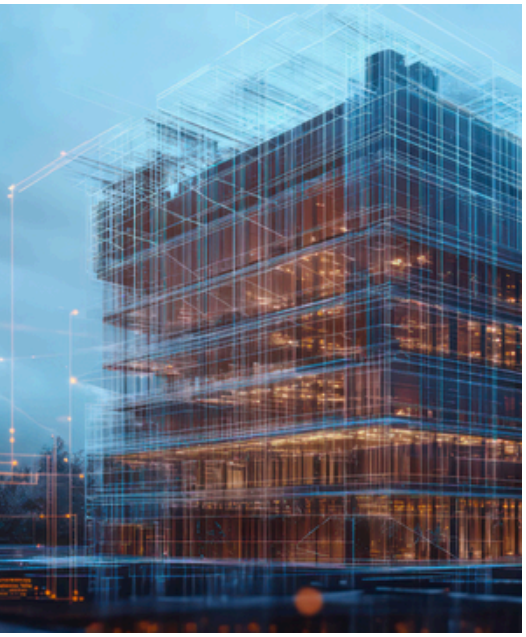

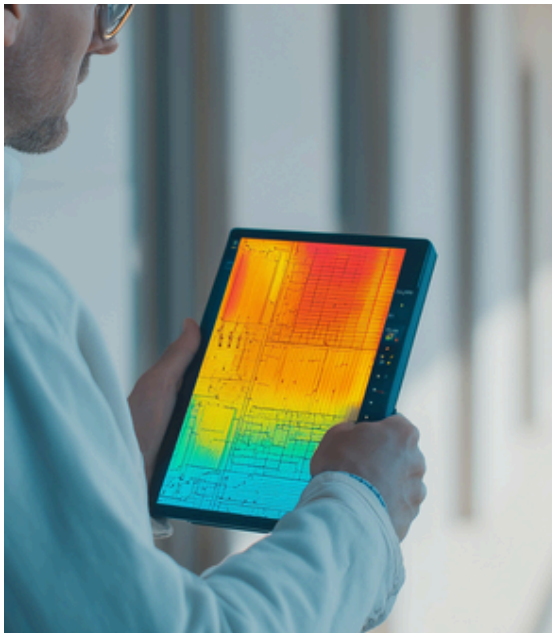

The Eng IndX and AgileRL Approach

Eng IndX and AgileRL partnered to build an HVAC control system that replaces static logic with an adaptive, learning-based architecture. The solution is grounded in Reinforcement Learning, a branch of AI in which an intelligent agent learns optimal behavior through structured interaction with its operating environment.

How it Works

The RL agent operates on a continuous feedback loop. It observes the building's current thermal state, including indoor temperatures, outdoor conditions, occupancy levels and HVAC system status. Based on this information, it selects a control action. The outcome of that action is evaluated against a **reward function** that balances two objectives: minimizing energy consumption and maintaining thermal comfort within defined bounds. Over time, the agent refines its control policy, learning the specific thermal dynamics of the building it operates in.

This approach differs fundamentally from rule-based systems. Rather than following predetermined responses to fixed thresholds, the agent develops a predictive understanding of how the building behaves. It accounts for thermal inertia, anticipates the effects of changing weather and adjusts proactively rather than reactively.

From Simulation to Production

The agent is initially trained in a simulated building environment, which allows rapid iteration without any risk to building operations or occupant comfort. Once baseline performance is established, the model is calibrated against real sensor data collected from the physical building. This calibration step is critical. It closes the gap between simulated behavior and actual thermal dynamics, ensuring that the control policy is robust and reliable before it goes live.

In production, the system operates in a fully closed-loop architecture. An orchestration layer integrates real-time sensor data, the trained RL agent and direct HVAC control interfaces. The agent continuously adapts, responding to changing conditions without manual intervention.

The Role of AgileRL in Accelerating Deployment

One of the most significant barriers to adopting RL in industrial settings is the engineering overhead required to move from a working prototype to a production system. Training orchestration, hyperparameter tuning, model versioning, deployment pipelines and performance monitoring all demand custom infrastructure that can take months to build. AgileRL addresses this directly. Its integrated RLOps paltform, Arena, provides **evolvable hyperparameter optimization**, automated training pipelines and end-to-end lifecycle management, all within a single, unified workflow.

For this project, AgileRL reduced hyperparameter optimization compute time by **40%** and eliminated the need for custom infrastructure development for model tracking and evaluation. This allowed the Eng IndX engineering team to focus on control performance and building-specific calibration rather than on operational tooling.

Results

Metric	Category	Detail
20%	**Energy Savings**	vs. conventional RL workflows
13%	**Comfort Stability**	within target temperature setpoints
30%	**Control Robustness**	reducing oscillations and mechanical stress
–35%	**Time to Deployment**	from development to production
–40%	**Compute Time**	for hyperparameter optimization

Beyond the headline metrics, the improved policy stability eliminated erratic switching behavior in the HVAC system. This reduction in mechanical cycling contributes to **smoother equipment operation and potentially extends hardware lifespan**, representing an additional maintenance cost benefit that compounds over time.

Business Impact

For building operators and real estate portfolio managers, the value of this solution is measurable and direct.

- **Lower energy costs:** A 20% improvement in energy efficiency translates directly to reduced utility spend, with payback periods that scale favorably across multi-building portfolios.
- I**mproved occupant experience:** More stable thermal conditions reduce complaint volume and support tenant satisfaction, a factor that directly influences lease renewal and occupancy rates in commercial real estate.
- **Reduced carbon footprint:** Energy reductions contribute measurably to Scope 1 and Scope 2 emissions targets, supporting ESG reporting and regulatory compliance.
- **Faster time to value:** The 35% reduction in deployment timelines means organizations see operational benefits sooner, with lower upfront integration costs.
- **Reduced engineering complexity:** AgileRL's Arena platform eliminates the need for bespoke ML infrastructure, making RL-based building control accessible to organizations without dedicated AI engineering teams.

Why Eng IndX

Eng IndX brings deep experience in deploying AI and data-driven systems in complex, operationally demanding industrial environments. Our partnership with AgileRL reflects a deliberate approach to building control: combining domain expertise in building operations and energy systems with production-grade reinforcement learning infrastructure.

This is not a research demonstration. It is a validated, deployable solution, already proven in real building environments and ready to scale across multi-site portfolios.

Next Steps

If your organization manages commercial building portfolios and is evaluating intelligent HVAC controls, we would welcome a conversation about how this solution applies to your environment. Contact Eng IndX at **info@indx.com** to schedule a technical briefing or pilot assessment.

About Eng IndX

Engineering Industries eXcellence is a division of the Engineering Group, a global leader in digital transformation. Focused on manufacturing and transportation, they deliver the Digital Thread for Industry 4.0 across the end-to-end lifecycle and supply chain of industrial products and processes. Learn more at **www.indx.com**.

About AgileRL

AgileRL is an RLOps company building Arena, an enterprise platform for developing and deploying reinforcement learning agents at production scale. Arena gives ML teams the infrastructure to train, evaluate and ship RL systems, with proven results including 2B parameter models reaching the performance of 8B models through evolutionary hyperparameter optimization. The company is headquartered in London with US operations in San Francisco. Learn more at **agilerl.com**.

About Decision Lab

Decision Lab is an award-winning decision intelligence consultancy. We are a team of experts in AI, optimization, and simulation dedicated to helping organizations make better decisions. Our solutions provide quantitative assurance for high-stakes capital investments and supply chain orchestration, with proven results across Defence, FMCG, Manufacturing, Pharma, and Utilities. Through our AI Innovation Lab, we pioneer the rigorous integration of simulation, mathematical optimization, and artificial intelligence technologies. The company is based in the UK. Learn more at **decisionlab.co.uk**.

Interested in speaking to one of our experts?

Contact us at indx.com | info@indx.com