

From energy consumer to flexible asset: the new role of data centres

12 June 2026, 13:01

Paul Peeters

Data centres that shift tasks and scale back consumption are changing the rules on the grid.

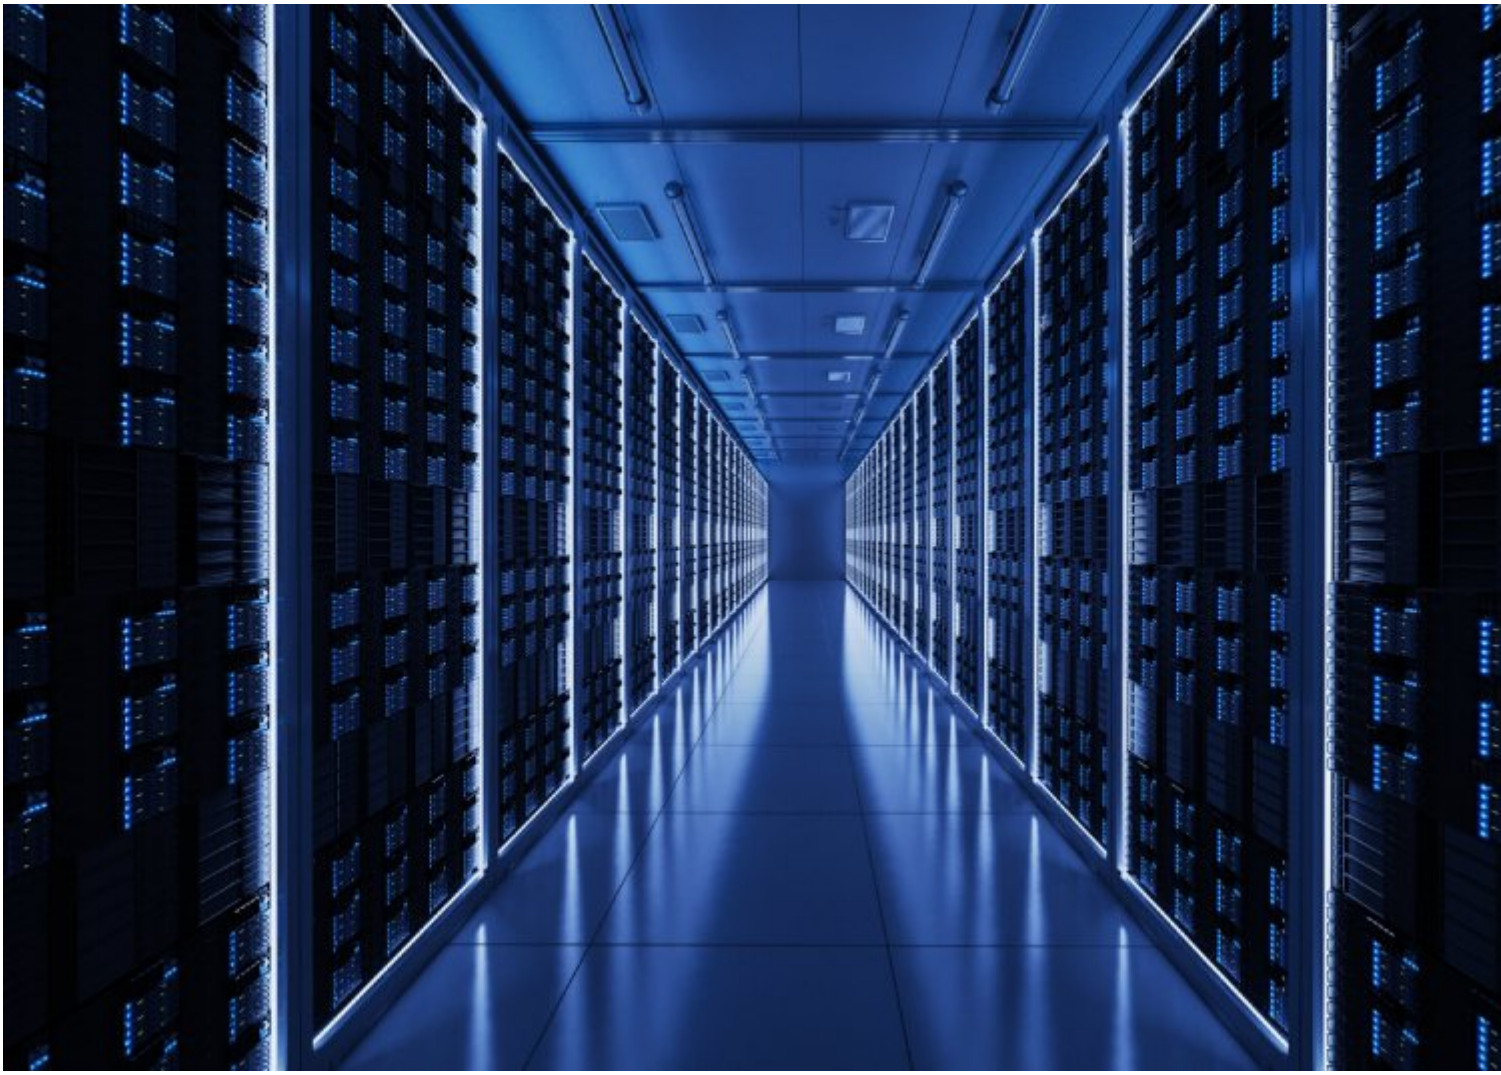
In a previous article, we looked at how companies are bypassing the queue to the grid by using behind-the-meter solutions and microgrids. But a more profound shift is taking place: the way in which data centres relate to the energy system is changing fundamentally. A data centre that can scale back its energy consumption, shift certain tasks to off-peak hours and supply reserves to the grid isn't a burden on the infrastructure. It's an asset. This shift in thinking has major implications for how we design, operate and regulate data centres.

The outdated view: the data centre as a passive consumer

The classic image of a data centre is that of a large, constant consumer of power: always on, always demanding, never giving anything back. That view is no longer accurate. Microsoft puts it bluntly: data centres are 'essential, not optimal'. They are neither a luxury nor a problem, but

essential infrastructure, comparable to the railways and electricity grids of the twentieth century. The challenge isn’t to make them smaller, but to use them more effectively.

The figures confirm this trend. Microsoft opened its first three data centres in Belgium in October 2025 and is investing ten billion dollars a month in data centre capacity worldwide. Migrating from an on-premise environment to Microsoft's hyperscale cloud delivers energy savings of up to 93 percent, and up to 98 percent lower carbon emissions. Revenue grew by 71 per cent, energy consumption by 160 per cent, but the efficiency gap is narrowing as new technologies gain traction.



Flexibility as a technical reality

Modern data centres are equipped with technologies that enable them to actively manage their energy consumption and thus integrate more effectively into the grid as a flexible load. This can be achieved through two strategies: geographical arbitrage (shifting the computing load to sites with cheaper or greener energy) and temporal arbitrage (scheduling tasks for times when prices are low or renewable energy availability is high). This enables data centres to participate in day-ahead, intraday, imbalance and ancillary services markets, and to respond to price and grid signals rather than simply accepting them passively:

- **Colocation**: distributing computing load across sites based on local energy prices and grid availability

- **Throttling**: temporarily reducing computing capacity during periods of peak prices or grid congestion
- **Timeshifting**: postponing non-urgent jobs to times when prices are low or there is high availability of renewable energy

At the hardware level, intelligent power distribution units (PDUs) provide an additional layer of optimisation. Automating dynamic load balancing across the phases every 24 hours can provide additional savings of 2 to 3 per cent. Servers that have been running unused for months are automatically detected and shut down.

Important context: most European data centres operate at 10 to 50 kilowatts per rack, not the 150 kilowatts or more used by hyperscalers. Flexibility techniques are therefore also relevant to mid-tier facilities, not just the major players.

The value of uptime: how flexibility is weighed up

Flexibility sounds appealing, but there is a tension: data centres are built with availability in mind. One hour of downtime costs an average data centre more than ten thousand dollars per megawatt-hour in lost revenue and SLA penalties. That is the benchmark against which every decision on flexibility is assessed.

That tension can be resolved, but it requires a nuanced approach. Not all computational processes are equally time-critical. Training processes for AI models can be rescheduled. So can rendering, batch processing and backups. Real-time inference and transaction processing cannot.

The business case for flexibility therefore depends on a thorough segmentation of the workload. Those who do this effectively combine multiple revenue streams from the same installation: optimising self-consumption, participating in the day-ahead market, intraday trading and ancillary services such as reserve capacity. Without this combination of revenue streams, the business case is rarely sound.



AI as an enabler of smart energy management

The flexibility described above is virtually impossible to manage without AI. The combination of price forecasts, weather data, demand and grid conditions is constantly changing. An energy management system that responds to this in real time is at the heart of smart data centre operations.

AI plays a dual role in this. Firstly, as a forecasting tool: more accurate forecasts of energy prices, peaks in demand and renewable energy production. Secondly, as a decision-making layer: which jobs should be carried out and when, which reserves should be made available, and when battery storage should be deployed. For this AI system to work, explainability is essential. Finance controllers must be able to verify energy bills and detect discrepancies. Once that trust has been established, experience shows that users stop monitoring granular details and let the automation do its job. Building that trust is the first step, not the last.

Fully unmanned data centres are already technologically feasible, but the sector is not yet ready for them. The human factor, and the associated issue of liability, is slowing down the adoption process. That will take some time.

What this requires of the sector

The technology for flexible data centres already exists. Operators' willingness is increasing. However, the sector faces a number of structural barriers that technical innovation alone cannot overcome. Data quality is the most underestimated problem. Energy management models are only as good as the data they are based on. Many systems produce incomplete, inconsistent or poorly labelled data. Anyone wishing to implement an AI-driven EMS would be well advised to start with a thorough audit of their data pipeline.

A related question that is being raised in the sector is: should we also introduce efficiency scores for AI models themselves? Hussain Kazmi (KU Leuven/EnergyVille) referred to a Hugging Face project that ranks models according to their energy consumption using a 1 to 5-star rating, similar to an EPC label for buildings. For the time being, this is a rough approach that doesn't adequately account for performance differences between models, but the idea is gaining ground.

Furthermore, the right stakeholders are still too often missing from the table. Distribution system operators (DSOs) and public authorities are essential for embedding data centre flexibility into the grid system on a long-term basis. The energy transition trilemma (security of supply, affordability, sustainability) and the data centre trilemma cannot be resolved independently. That requires more than just technology: it requires shared frameworks, regulation and investment certainty. For companies looking to get started today, a pragmatic approach is needed: first optimise existing processes, then build predictive models and only automate what has been shown to work. The first successful application will fund the next step.

How can you make your data centre or industrial energy system flexible?

Sirris informs and advises companies about implementing smart energy management, EMS, and flexibility strategies, from initial analysis through to operational implementation.

Get advice from our experts

Discover the other parts of AI 4 ENERGY.

Part 1: The queue for the grid: why behind-the-meter is becoming strategic
Part 2: From energy consumer to flexible asset: the new role of data centres ?

Authors



Paul Peeters