

B E N C H M A R K I N G R E P O R T



WHEN AI BRAKES THE MODEL

The New Economics Of Data Center
Opex

Foreword

October 2025

AI has broken the long-standing cost model of data centers, shifting stress points and pushing cost, energy, and design to the center of strategic decision-making.

For over a decade, Atomiton has helped organizations measure and manage their carbon and energy performance, and this report reflects both our market research and our experience supporting operators navigating the shift to AI-dense compute.

My hope is that the report not only highlights the rising cost pressures, but also equips you with practical levers to act. The coming years will reward operators who embrace efficiency, strategic siting, and AI-ready infrastructure design.



Jane Ren
Chief Executive Officer
Atomiton, Inc.

“AI demands more than capacity expansion. It demands foresight in energy procurement, efficiency, and siting. The winners will be those who act before cost pressures force their hand.”

Executive Summary

The economics of data centers have entered a new phase. For more than a decade, data centers benefited from falling unit costs: efficiency gains, virtualization, and economies of scale steadily reduced Opex per kW. That era is over. In the last two years, global volatility in energy markets, inflationary pressures, and the explosive growth of AI workloads have combined to reverse the trend. Operators must now adapt to a reality where the cost to operate each kilowatt of capacity is rising, not falling.

Energy is no longer a background line item — it is now strategic.

Electricity has shifted from being a predictable operating expense to a board-level risk. It is simultaneously the largest driver of cost increases and the most visible component of carbon reporting. Operators cannot treat energy as a pass-through; it is a strategic variable that defines competitiveness, location advantage, and long-term viability.

AI workloads are reshaping the unit economics of compute. Traditional workloads allowed operators to amortize costs across moderate rack densities and low utilization. AI turns that model upside down: racks 5–10× more power-hungry, running at near-full load, with cooling demands that overwhelm legacy infrastructure. For those unprepared, AI is not just a growth opportunity but an accelerant of cost and risk.



The efficiency gap between leaders and laggards is widening.

Hyperscalers are sustaining Opex per kW at half the level of enterprises, leveraging scale, cheap power, and purpose-built designs. Colocation providers sit in the middle, but even here, those that adapt fastest to high-density AI will capture share. The result is a bifurcating market: operators with foresight pull further ahead, while underutilized or outdated sites risk becoming stranded assets.

Strategic focus must shift from cost control to resilience and positioning.

The challenge is no longer squeezing incremental savings. The new imperative is resilience — securing stable, low-cost energy, designing AI-ready infrastructure, and embedding efficiency as a strategic differentiator. Operators that treat Opex per kW not just as a cost metric, but as a competitive weapon, will define the next decade of digital infrastructure leadership.



Benchmarking Insight


In this report, all cost per kW or cost per kW-year figures are normalized against the commissioned IT kW (installed IT capacity) of the facility, not the utility power feed or the mechanical/electrical headroom.

That means for the metrics: Numerator = annualized Opex (staffing, M&E maintenance, energy where included). Denominator = commissioned IT kW (the amount of IT load the site is designed and equipped to support).

1. Opex per kW fell for a decade, but has risen ~10–20% since 2022

For most of the 2010s, Opex per installed kW declined steadily, reflecting efficiency gains, consolidation, and economies of scale. That trajectory reversed in 2022: inflation, higher energy tariffs, and supply-chain pressures pushed unit Opex upward across all data center types. By 2024, hyperscale and colocation operators are spending ~10–20% more per kW than in 2021, while enterprises, more exposed to labor and inefficiency, have seen even sharper increases. The era of “ever-cheaper” Opex per kW is over; operators must now plan for structurally higher costs


Opex per kW Trending



2. Energy cost per kW spiked in 2022 and remains 30–50% above pre-crisis baseline

Electricity costs surged worldwide in 2022, driven by commodity shocks and grid volatility. While prices moderated somewhat in 2023–24, the median energy cost per kW of IT load remains 30–50% above the pre-crisis baseline. This creates lasting upward pressure on total Opex, particularly for colocation and enterprise facilities that cannot hedge or negotiate rates at hyperscale. Energy volatility has become a structural feature of data center operations, demanding active procurement strategies.


Global Energy Cost (USD) per kW Trending



3. Hyperscalers operate at ~½ the Opex/kW of enterprise facilities (2024)

In 2024, hyperscale operators achieve Opex per kW of ~\$3,400 versus ~\$10,300 for enterprise sites — nearly a 3× efficiency gap. Scale, automation, and optimized designs allow hyperscalers to spread fixed costs and secure cheaper energy, while enterprises suffer from smaller footprints, lower utilization, and higher redundancy. Colocation sits in the middle at ~\$5,600/kW. This efficiency gap underlines why enterprises increasingly turn to cloud or colo for power-intensive workloads.


Opex per kW vs. Data Center Type (2024)



3. Energy share of Opex ranges from ~20% (hyperscale) to ~50% (enterprise) in 2024

Energy's share of Opex reveals structural efficiency: hyperscalers keep energy at ~20–25% through low tariffs and lean staffing, while colocation providers average ~40%, and enterprises often exceed 50%. A high energy share indicates vulnerability to price shocks and efficiency drag. Enterprises that allow energy to dominate their cost base risk margin erosion; targeted PUE improvements and smarter power contracts can rebalance the cost mix.


Energy Share of Opex vs. Data Center Type (2024)



5. Colo economics: clients pay fees + metered energy; operators carry ~\$4–5k/kW facility Opex (2024)

Retail clients typically pay ~\$2,600 per kW-year in recurring fees, wholesale clients ~\$1,800 — excluding metered energy, which averages ~\$1,600. Behind the scenes, operators carry ~\$4,500 per kW in facility Opex, covered by the markup between costs and recurring charges. Rising energy prices squeeze this margin and force colos to increase lease rates, as seen in major markets over the past 24 months. Operators must strike a balance: recover higher costs without eroding customer demand.


Colo Fees per kW year vs. Operator Opex (2024)



6. AI racks draw 5–10× more power than traditional racks (2023 - 2024)

Traditional workloads consume 5–10 kW per rack; AI racks often exceed 50 kW. This magnitude shift redefines the economics of cooling, power distribution, and space utilization. A single AI rack can equal the load of half a legacy row. Facilities designed for 10 kW racks face disproportionate cost and risk when accommodating AI gear, making proactive high-density design essential for competitiveness..


Typical Rack kW vs. Profile (2024)



7. Hybrid facilities see disproportionate cooling overhead for small AI clusters

When a small number of AI racks are retrofitted into an air-cooled environment, the entire facility's cooling efficiency suffers. Overhead power consumption rises sharply (equivalent to PUE ~1.9 vs ~1.6 baseline), inflating energy costs per kW. By contrast, purpose-built liquid cooling can contain overhead to PUE ~1.3. Operators who delay dedicated AI cooling investment risk higher costs across the board — penalizing both AI and traditional workloads.


Cooling Overhead kWh per kW year vs. Scenario



8. Energy cost per kW varies 3–5× by geography

Geography remains the single strongest driver of energy cost per kW. In 2024, Nordic sites average ~\$800/kW-year versus ~\$1,300 in the US, ~\$2,000 in the EU, and as high as ~\$3,500 in Europe’s 2022 crisis peak*. APAC markets vary widely but often run above \$2,000. This spread underlines the strategic value of siting AI-dense workloads in low-cost, renewable-rich regions. Operators in high-cost markets must lean on efficiency and contracts to remain competitive.

Energy Cost (USD) per kW year by Region



** In 2022, Europe experienced an unprecedented spike in electricity prices, driven by the energy crisis following the Russia–Ukraine war. Gas supply disruptions and market volatility pushed wholesale power prices in many EU markets above €300/MWh, with some spot prices surging even higher. For data centers, this translated into effective energy costs of ~\$3,500/kW-year — nearly 4× the long-term European average.*


We include this “EU 2022 peak” bar in the chart not as a current benchmark, but as a stress-test reference: it demonstrates how exposed operators in high-cost, fossil-dependent markets are to geopolitical and commodity shocks.

The “EU average” bar reflects normalized run-rates for 2023–24 (around ~\$2,000/kW-year). It does not permanently bake in the 2022 peak, which was a temporary spike. Instead, it represents a “stabilized” post-crisis average — still above pre-2021 levels, but far below the 2022 shock.

9. Unit cost gap: underutilized facilities pay >2× per active kW (2024)

Analysis of 2023–24 benchmark data shows that installed Opex per commissioned IT kW typically averages around \$5,000. When utilization diverges, the effective cost per active kW shifts dramatically: at 80% utilization, the cost is about \$6,250; at 35%, it jumps to over \$14,000 — more than double. These figures are modeled directly from industry-reported Opex benchmarks and utilization ranges (Uptime, CBRE, Ponemon), making the example representative. The implication is clear: underutilization is the silent driver of high enterprise costs, often outweighing gains from efficiency upgrades. The most powerful lever to reduce cost per active kW is consolidation — shutting under-filled sites or densifying workloads before adding new capacity.


Range of Annual Opex (USD) per Active kW



10. Energy % of Opex peaked in 2022; efficiency investments are bringing it down

Energy rose to 52% of Opex at the height of the 2022 energy crisis. Since then, investments in cooling upgrades and efficiency have pulled it back to ~47%. Still above the pre-2021 ~38%, energy remains the largest controllable lever in Opex. The clear takeaway: energy efficiency programs deliver measurable cost relief and should remain a board-level priority as AI drives absolute energy use higher.


Energy Cost as Share of Opex (%) Trending



Six Moves to Win in the Age of AI Economics


The challenge is clear: cost control alone is no longer sufficient. Part 3 introduces six strategic moves that operators can take to secure resilience, unlock competitiveness, and lead in the new economics of digital infrastructure.

Move 1: Secure long-term power contracts in low-cost regions




The challenge is clear: cost control alone is no longer sufficient. Part 3 introduces six strategic moves that operators can take to secure resilience, unlock competitiveness, and lead in the new economics of digital infrastructure.

Move 2: Benchmark and optimize PUE ahead of AI deployments



AI workloads magnify inefficiencies: a PUE of 1.8 vs. 1.3 can double the cooling cost burden at 50 kW racks. Before introducing AI gear, operators should benchmark current PUE, identify airflow or cooling inefficiencies, and implement low-capex improvements (e.g., hot aisle containment, raised temperature setpoints). Entering AI densification with subpar PUE exposes operators to runaway costs and capacity constraints. Target <1.5 PUE as a baseline for any facility planning to host AI racks.

Move 3: Reallocate AI workloads to facilities with lowest \$/kW energy



Not all data centers are equal. Energy cost per kW varies up to 5× by geography. Allocating AI clusters to low-cost regions (Nordics, Midwest U.S., hydro-rich APAC) can reduce effective Opex by thousands per rack per year. This requires portfolio-level planning: treating workloads as mobile assets, matching them with cost-optimal facilities. Enterprises should consider hybrid strategies — keeping latency-sensitive apps local, but pushing AI training to low-cost, high-efficiency zones.

Move 4: Plan liquid cooling retrofits proactively, not reactively



AI racks cannot be sustained with traditional air cooling beyond ~20 kW. Retrofitting liquid cooling under duress is costlier and riskier than proactive design. Operators should pilot small-scale deployments, establish vendor partnerships, and budget for gradual rollout. Proactive planning cuts retrofit costs by 20–30% and avoids disruption to legacy workloads. More importantly, it positions operators as “AI-ready” in the eyes of customers and investors.

Move 5: Track energy cost per kW as a lead KPI, not a lag metric



Most operators monitor PUE but fail to track \$/kW energy as a management KPI. Treating it as a lead metric allows executives to anticipate Opex creep before it hits budgets. Embedding this KPI in quarterly reviews forces accountability: if energy cost per kW breaches the 75th percentile, immediate corrective actions (renegotiating tariffs, efficiency retrofits, workload moves) can be triggered. The discipline turns an external cost risk into a controllable lever.

Move 6: Deploy analytics to optimize compute energy use (optional but high leverage)



AI and hybrid workloads introduce complex energy-use profiles. Analytics and workload-aware orchestration tools can optimize when and where jobs run — shifting batch AI training to off-peak hours, aligning workloads with renewable availability, or consolidating compute to minimize idle draw. For enterprises and colos, investing in such analytics can reduce effective energy spend by 10–20% without new hardware. Early adopters position themselves to manage not just cost, but also carbon reporting obligations.



Appendix A. References

1. Uptime Institute – Global Data Center Survey Reports, 2022–2024.
2. AFCOM – State of the Data Center Industry Reports, 2023–2024.
3. CBRE – Global Data Center Market Trends, 2024–2025.
4. Cushman & Wakefield – Global Data Center Market Comparison, 2023.
5. Ponemon Institute / Emerson Network Power – Cost of Data Center Outages, 2016; Operational Benchmarking Studies.
6. Datacenters.com – Data Center Cost Analysis, 2024.
7. Thunder Said Energy – Hyperscale Data Center Economics, 2024.
8. DOE / Lawrence Berkeley National Laboratory (LBNL) – Data Center Energy Efficiency Research, 2018–2023.
9. Vertiv – Data Center Energy Outlook, 2023.
10. Industry Case Studies & Vendor Documentation (NVIDIA, Dell, HPE, Enconnex, etc.), 2023–2024.