

The Power Equation

Measuring What
Really Matters in
AI Infrastructure

AIRSYS
Balance the Environment



Contents

- Introduction: The Efficiency Paradox 4
- Efficiency Vs. Effectiveness:
A New Equation for
Understanding Power Value. 6
- Measuring What Matters:
Introducing PCE and ROIP 8
- AIRSYS: Redefining What
Cooling Can Enable 10
- Conclusion: The Power is in Your Hands 14

Introduction

The Efficiency Paradox

“At current rates of AI workload growth, Virginia will not produce enough electricity to support unconstrained data center growth.”

This is the troubling trajectory the industry is accelerating toward, despite years of steady efficiency gains. For nearly two decades, the industry has chased better PUE, with average values improving from roughly 2.5 in 2007 to 1.56 in 2024, according to Uptime Institute’s 2024 Global Data Center Survey.

But even with these improvements, operators are still running up against a fundamental limit: power.

The explosive scale of AI is what ultimately turns power into the most critical and constrained resource for data centers. Deloitte projects that U.S. AI data center power demand will surge from 4 GW in 2024 to 123 GW by 2035, a rise so steep that it will outstrip available power supply. This scenario is outlined in a December 2024 report (source of the quote above) submitted to the Governor and General Assembly of Virginia — the world’s largest data center market.

This growing mismatch between demand and available power creates the industry’s efficiency paradox: **data centers have never been more efficient, yet they have never been more constrained.** Even with excellent PUEs, facilities still operate within hard power limits while delivering ever-growing amounts of compute. Power — not cooling, not space — has become the defining bottleneck of the AI era.

This is where the industry’s conversation changes. The question is no longer “How efficient are we?” but “How much compute are we getting out of the power we already have?”

This brief introduces two metrics that complement PUE and address the power-to-compute challenge. Armed with these insights, data center operators and investors can better understand the true value of their power allocations and make more informed decisions about capacity planning and operational effectiveness.



Efficiency vs. Effectiveness

A New Equation for Understanding Power Value

Efficiency reflects how little energy you waste, but do you know how much output you actually generate? It's sometimes easy to forget that energy efficiency does not equal power productivity, and this visibility gap carries significant operational, strategic, and financial implications:

“WASTED” REDUNDANT POWER

Redundant capacity is necessary for data centers to maintain reliability and support scalability. Yet, a significant portion of the power reserved for N+1/N+2 reliability rarely, if ever, feeds IT loads. Instead, it sits idle within the power chain, consuming power without contributing to compute production. As a result, it masks true power productivity by making capacity look usable when it's ultimately unavailable.

“Up to 30% of all servers in any given data center may be ‘zombies’ = servers that consume energy but serve no useful purpose” – Lawrence Berkeley National Laboratory

STRANDED POWER = LOST OPPORTUNITIES

Beyond redundancy-related waste, even active servers rarely operate at more than 50% utilization. When power availability is already constrained, this underutilized capacity becomes a lost opportunity, resulting in stranded power that translates into lost revenue. One CompuDynamics case study found that a campus with 5 MW of unused contracted power was effectively losing nearly \$10 million per year. In other words, stranded power is the silent killer of compute growth, and the returns it should generate.

“An estimated 10 million servers sit completely idle, representing \$30 billion in wasted capital.” – Fortune

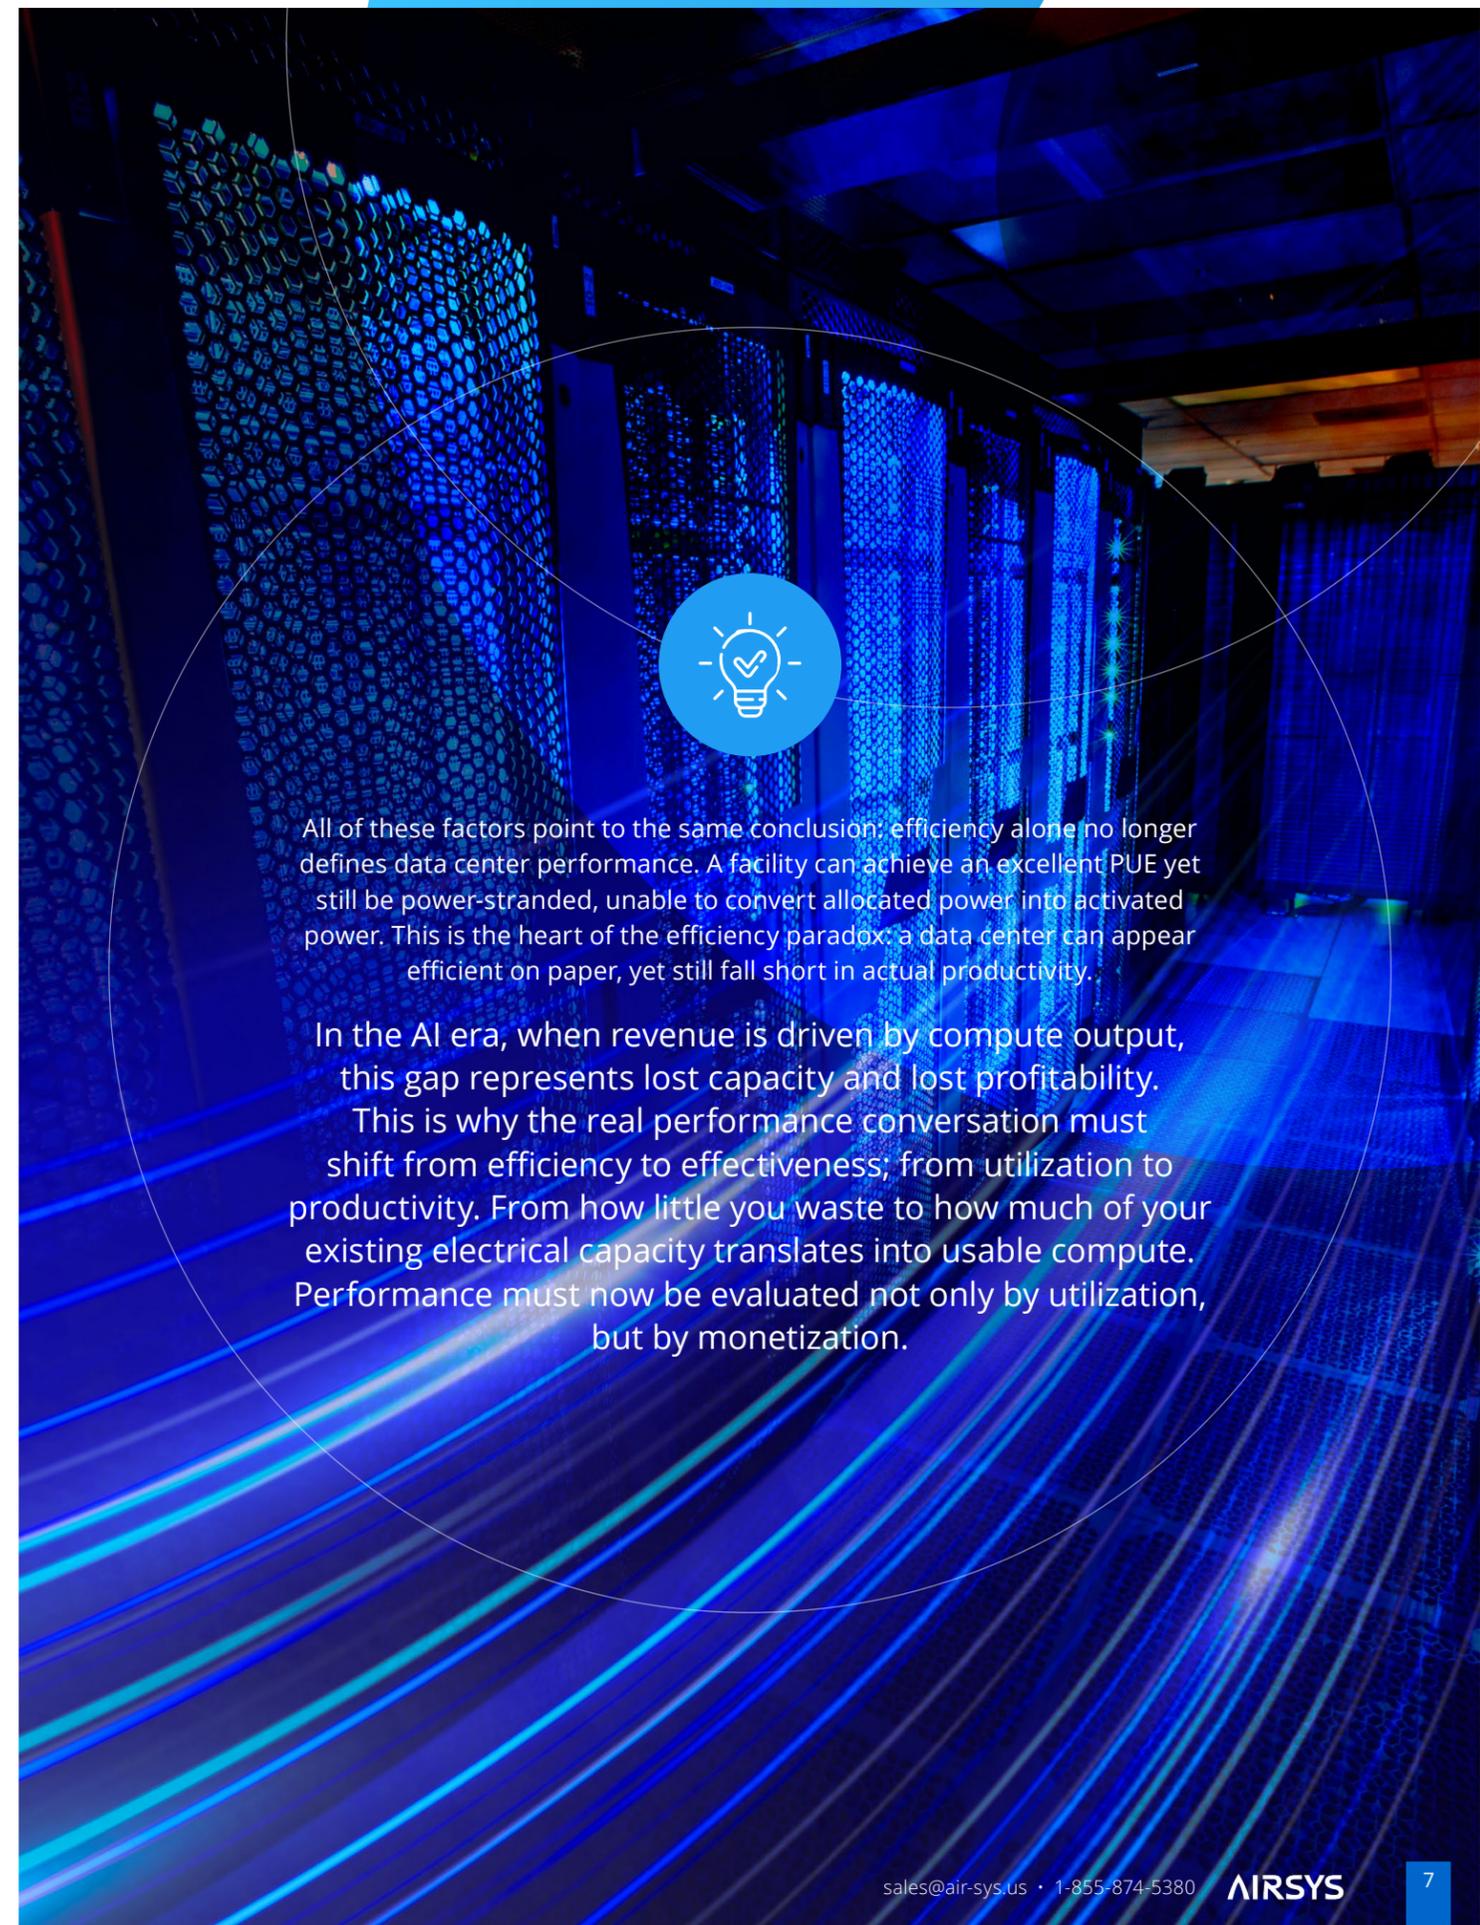
POWER AVAILABILITY DICTATES DATA CENTER PROJECTS

Data center construction has historically been driven by factors such as building costs, tax incentives, and geographic risk, but the rise of AI has pushed power availability to the top of the list. Build projects are increasingly delayed, downsized, relocated, or canceled based on power availability and utilization effectiveness. In constrained markets such as Virginia, the Netherlands, and Ireland, project approvals can even stretch to 10 years. This puts expansion plans, revenue timelines, and market-entry strategies on the line.

“Data center construction completion timelines have been extended by 24 to 72 months due to power supply delays.” – CBRE

BLIND SPOTS IN PERFORMANCE EVALUATION

When PUE improves, data center stakeholders often assume operations are strong. But as energy becomes the new focus for data center growth, they move from traditional efficiency metrics to asking a more fundamental question: **how much power can be made available, fast enough and affordably enough, to support growth?** Without a clearer view of productive power, decisions are made on perceived progress rather than real operational capability.



All of these factors point to the same conclusion: efficiency alone no longer defines data center performance. A facility can achieve an excellent PUE yet still be power-stranded, unable to convert allocated power into activated power. This is the heart of the efficiency paradox: a data center can appear efficient on paper, yet still fall short in actual productivity.

In the AI era, when revenue is driven by compute output, this gap represents lost capacity and lost profitability. This is why the real performance conversation must shift from efficiency to effectiveness; from utilization to productivity. From how little you waste to how much of your existing electrical capacity translates into usable compute. Performance must now be evaluated not only by utilization, but by monetization.

Measuring What Matters: Introducing PCE and ROIP

Power, Productivity, and Profitability in the AI Era

If power efficiency alone no longer defines performance, the industry needs new indicators that illuminate what efficiency fails to show: how much of a facility's megawatts become meaningful compute, and how much value those megawatts ultimately generate. To answer these questions, we introduce two new metrics that enhance power performance evaluation: **PCE** (Power Compute Effectiveness) and **ROIP** (Return on Invested Power).

Together, they link power, compute, and capital efficiency, giving operators, investors, and partners a direct view of how well electrical capacity is monetized and how effectively each kilowatt contributes to return on investment and usable intelligence.

POWER COMPUTE EFFECTIVENESS (PCE): HOW MUCH POWER BECOMES COMPUTE

PCE reveals how much of your provisioned power becomes real compute, surfacing idle or underutilized capacity.

- » Formula: $\text{IT Load (kW)} / \text{Total Power Capacity (kW provisioned)}$
- » Direction: Higher is better
- » Range: 0.50 (inefficient) to 0.90+ (highly effective)

RETURN ON INVESTED POWER (ROIP): WHAT YOU EARN FROM EVERY KILOWATT

ROIP shows the financial return on each kilowatt invested. High ROIP means power is producing revenue-generating compute, while low ROIP indicates megawatts are being consumed with little economic output.

- » Formula: $\text{Revenue or Compute Value (\$ or TFLOPs)} / \text{Power Invested (kW) (total active power used)}$
- » Direction: Higher is better

WHY PCE AND ROIP MATTER

Elevate Power's Strategic Importance

- » Reframe revenue expectations around compute output rather than physical footprint.
- » Resolve growth bottlenecks created by grid and permitting limitations.
- » Integrate sustainability goals with PCE and ROIP that complement carbon-intensity and water-use metrics.

Expose Hidden Capacity and Improve Performance

- » Identify the power that is provisioned but not in use.
- » Convert hidden capacity into activated compute.
- » Redirect power consumed by cooling into compute output, where it generates greater value.

Strengthen Financial and Investment Decisions

- » Quantify the revenue implications of underutilized power.
- » Evaluate true growth potential before investing.
- » Justify capacity expansions or optimization projects.
- » Connect power effectiveness directly to ROI and ESG performance.

HOW THE THREE METRICS WORK TOGETHER

It's important to highlight that PCE and ROIP do not replace PUE. All three metrics must be viewed together to see the complete picture of power performance and create a stronger foundation for operational, financial, and expansion decisions.

METRIC	MEASURES	DIRECTION	USE	INTERPRETATION
PUE	Energy overhead	Lower is better	Cooling efficiency	How efficient the facility is
PCE	IT utilization (technical effectiveness)	Higher is better	Capacity utilization	How productive the facility is
ROIP	Power ROI (economic effectiveness)	Higher is better	Financial productivity	How profitable the facility is



AIRSYS

Redefining What Cooling Can Enable



For more than 30 years, AIRSYS has redefined heat management in mission-critical environments, approaching cooling as a strategic driver of performance, rather than a functional necessity.

Now, as power constraints tighten and AI workloads push thermal demands to new extremes, AIRSYS responds with efficient, scalable, and sustainable solutions that recover capacity, enhance resilience, and prepare facilities for the next wave of high-density growth.

PowerOne™

Cooling Built for the AI Century

Cooling already accounts for up to 40% of total data center power use, yet every kilowatt spent on cooling is a kilowatt not driving compute. PowerOne is built to change that.

PowerOne is AIRSYS' flagship portfolio of intelligent, high-efficiency cooling systems. Designed to empower ultra-high-density IT loads, PowerOne acts as an enabler of higher compute output, freeing up power lost to inefficient cooling to drive more computing. In an energy-constrained world, PowerOne sits at the center of the utilization-to-productivity transition, allowing operators to convert more of their existing power into usable, revenue-producing IT load.

Traditional cooling improves efficiency. PowerOne restores capability.



LiquidRack™

The Next Evolution of Intelligent Liquid Cooling

Data centers are entering an era where rack densities reach hundreds of kilowatts, and legacy cooling methods are hitting their limits. LiquidRack is AIRSYS' solution — the next-generation of server-level liquid cooling technology designed specifically for the thermal profile of modern AI and GPU-dense architectures.

HOW DOES IT WORK?

LiquidRack uses patented single-phase spray cooling, which applies dielectric coolant directly to CPU and GPU surfaces for immediate heat removal. The system operates through two separate cooling circuits (a coolant loop and a water loop) connected by a high-efficiency heat exchanger.

1. Cold dielectric coolant is sprayed onto waterproof servers to absorb heat at the source.
2. Heated coolant returns to the plate heat exchanger, where the heat is transferred into the water circuit.
3. Heated water is then rejected to the atmosphere via a dry cooler, completing the heat-removal cycle.

By eliminating heat directly at the chip, LiquidRack delivers faster heat rejection, tighter temperature control, and significantly improved thermal stability. This higher cooling efficiency means more of the facility's power can be directed from wasteful cooling toward computing production.

BREAKTHROUGH ADVANTAGES

- » Uses 80% less coolant in an 80% smaller footprint
- » Supports input capacities up to 170 kW per rack
- » Reduces system-level cooling infrastructure
- » Enables 100% energy recovery from IT load
- » Deployable in both new builds and retrofits
- » Built to scale with future processors and AI training clusters



Cloud Service



Data Centers



Crypto-mining



Telecom Facilities



AI Technology

CASE STUDY: TRANSFORMING A LEGACY FACILITY WITH LIQUIDRACK

A 10 MW legacy data center with mechanical cooling dedicated 5–6 MW to cooling and 1–2 MW to infrastructure, limiting IT load to approximately 10MW.

After retrofitting with LiquidRack spray cooling:

- » Cooling load dropped to 1 MW
- » IT load increased to ~15 MW
- » PCE improved from 0.63 to 0.93
- » ROIP increased by more than 35% due to higher compute revenue per kW

LiquidRack signals AIRSYS' progression in the direction the industry is headed, turning cooling into a catalyst for greater compute output. This is a first demonstration of what innovative, efficient, high-density cooling can achieve.

Click or scan the code below to visit our website to learn more about LiquidRack and download the datasheet.



Conclusion

The Power Is in Your Hands

Across the industry, a new reality has taken hold: power literally has the power to determine which data center projects move forward and which are forced to stall. Virginia's power constraints serve as a proof point, showing that even the world's largest data center market can outgrow its power supply.

This makes one thing clear: power efficiency alone is no longer enough to guarantee progress.

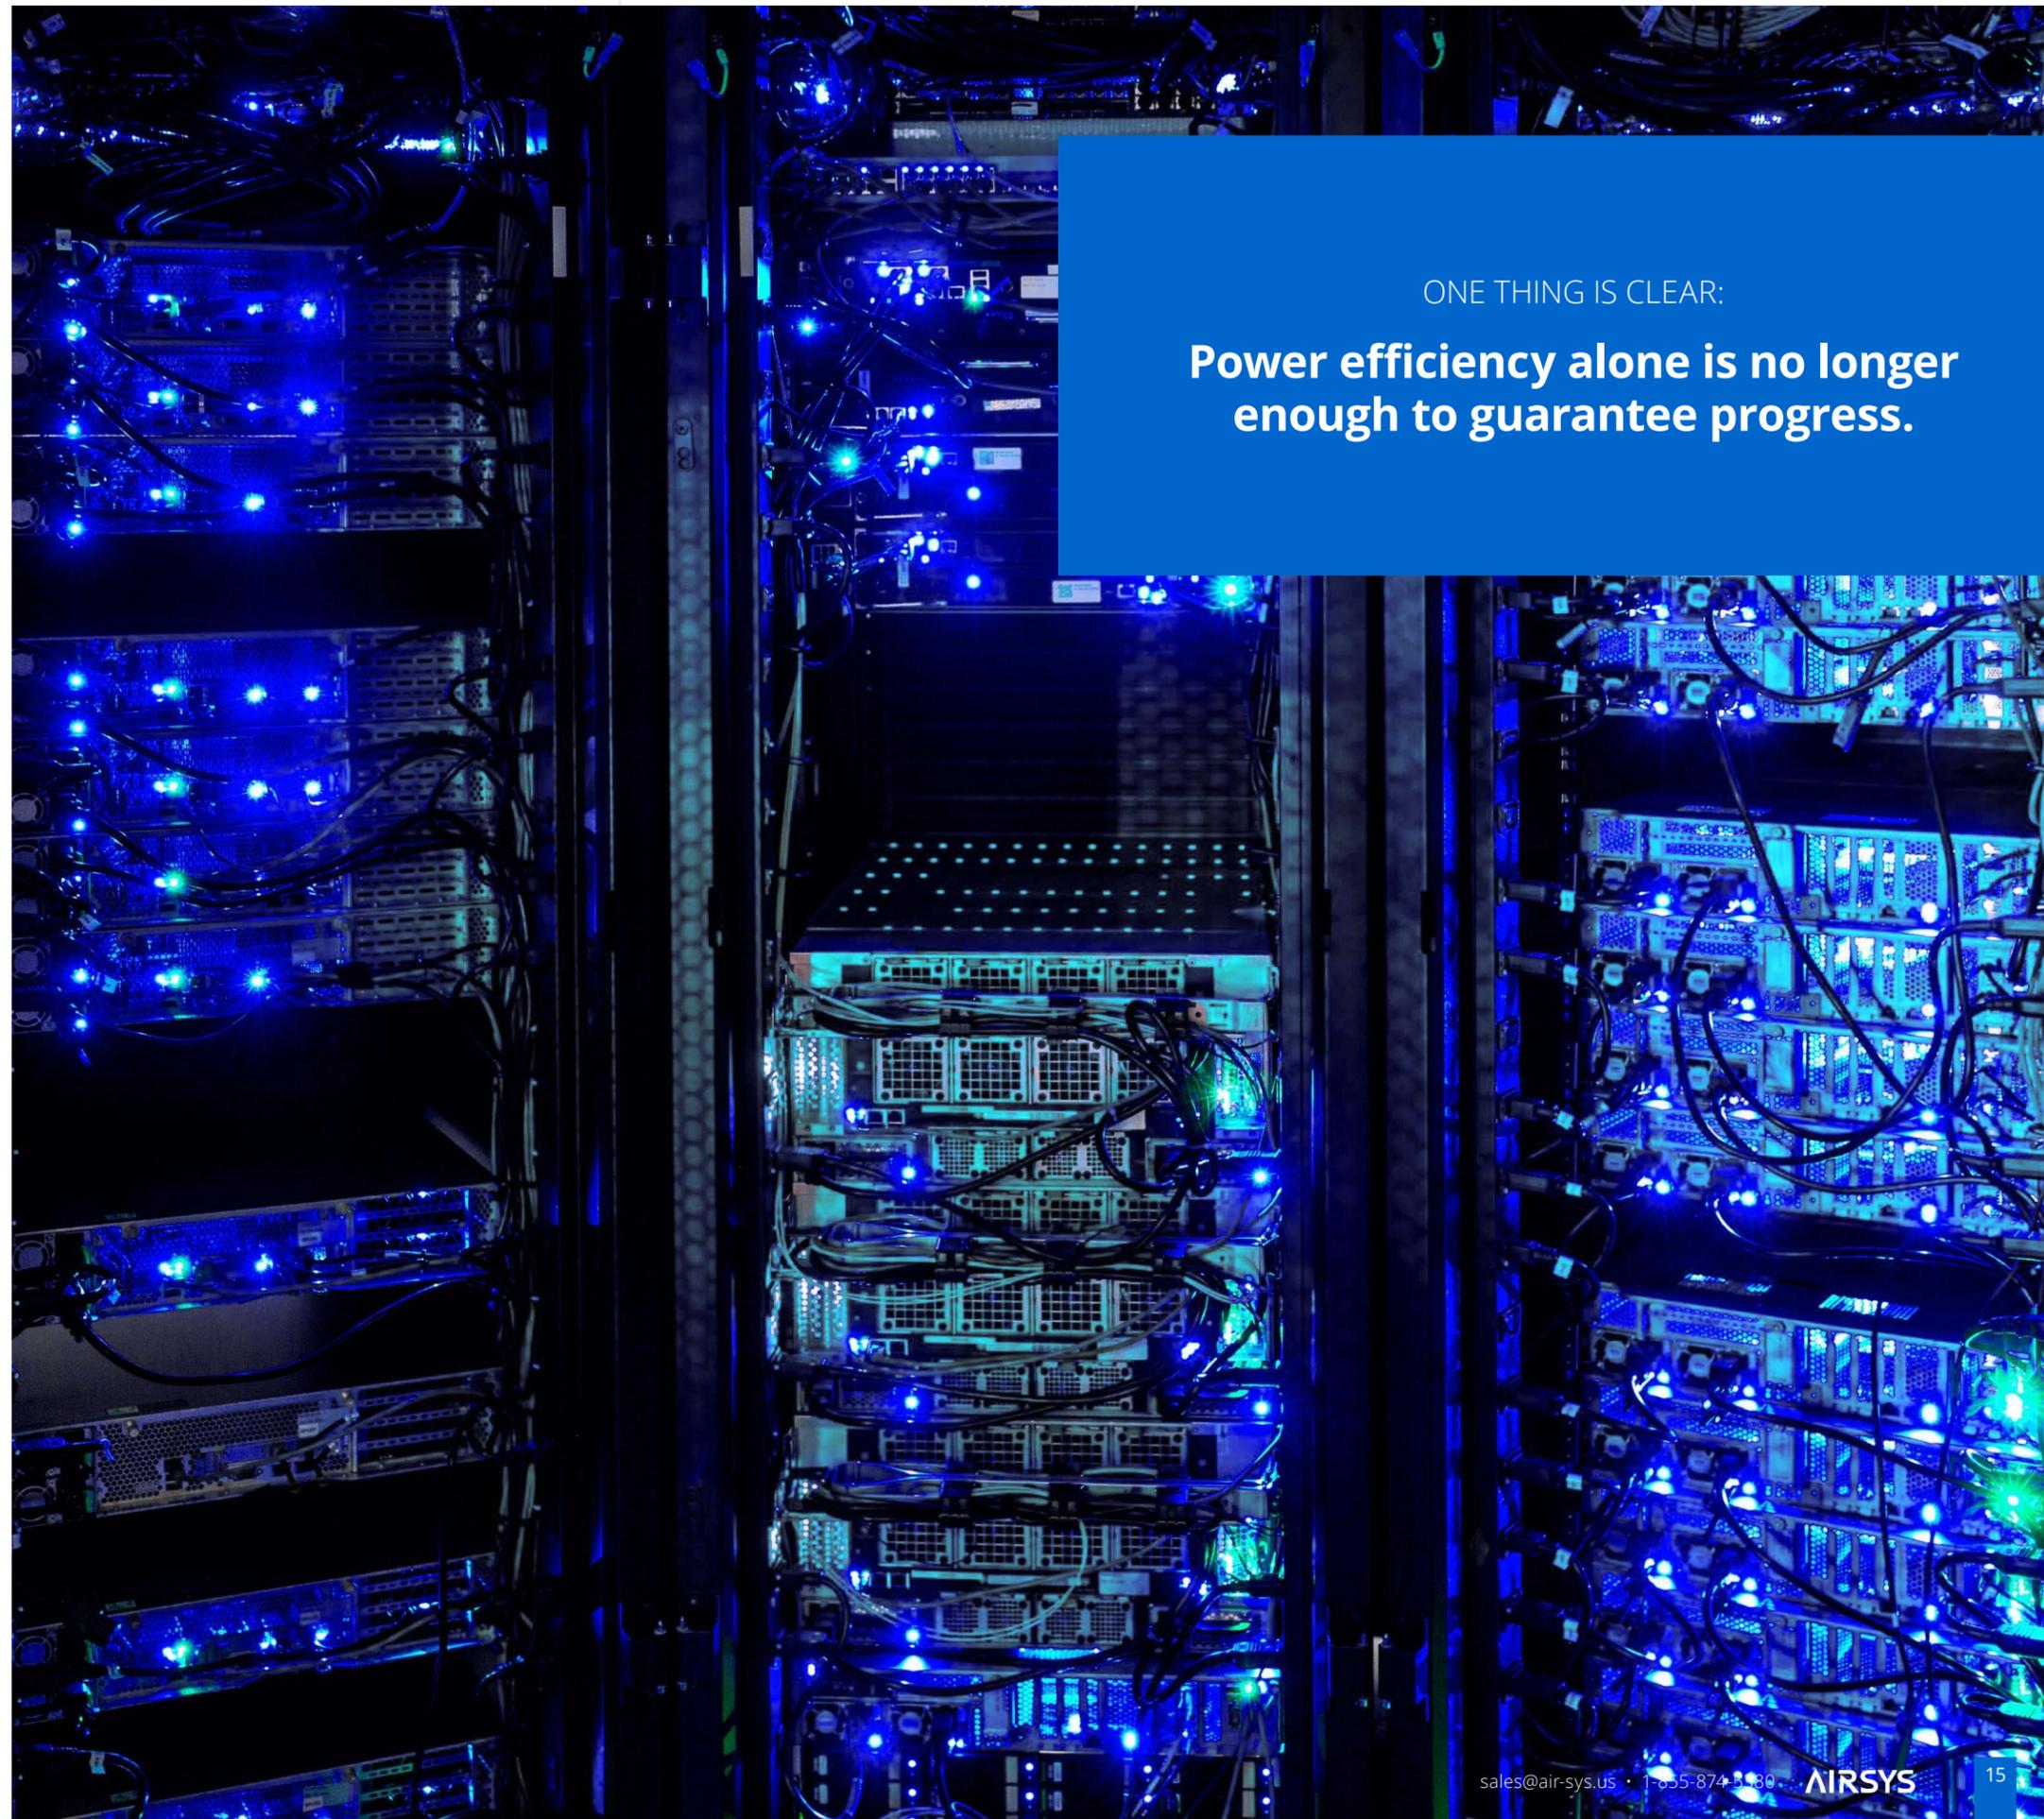
To navigate confidently in this landscape, operators and investors should begin focusing on a new set of priorities:

- **Optimize for productive power:** Measure how much capacity actually turns into compute.
- **Reduce the sources of stranded power:** Idle servers, oversized cooling, and unnecessary electrical overhead directly limit revenue potential.
- **Treat cooling as a strategic enabler:** The right cooling architecture can expand usable IT load without waiting on new utility power.
- **Design for scalability under constraints:** Power availability may not grow at the pace of demand, but compute effectiveness absolutely must.
- **Adopt metrics that reflect real capability:** Tools like PCE and ROIP separate perceived progress from actual operational performance.

The next generation of AI factory data centers will not be defined by who has the most power, but by who uses their power best. Those who shift from efficiency thinking to effectiveness thinking, from counting kilowatts to capturing compute, will lead the market into its next phase of growth.

**So, where will you take your facility from here?
The power is in your hands.**

ONE THING IS CLEAR:
**Power efficiency alone is no longer
enough to guarantee progress.**



Let's Talk!

Bridge the gap between today's infrastructure and tomorrow's AI-fueled thermal challenges with a system built for resilience, scalability, and performance.

Together, we'll keep your data center ready for what's coming.



Scan to request a FREE consultation or give us a call at 855-874-5380 to learn more.

AIRSYS
Balance the Environment

7820 Reidville Road, Suite 100
Greer, South Carolina 29651, USA

1-855-874-5380

airsysnorthamerica.com