

# Optimized Power Management for Air-Cooled Data Centers

*Toward energy-per-workload optimization in air-cooled data centers*

A technical white paper prepared for ASHRAE TC 9.9 committee review

April 2026

<b>Prepared for</b>	ASHRAE Technical Committee 9.9 audience
<b>Subject</b>	Closed-loop control of thermal headroom, airflow, and workload energy productivity
<b>Source basis</b>	Consolidated from the two provided OPM technical white paper drafts
<b>Document type</b>	Standards-style technical white paper consolidated draft

## Executive Summary

This consolidated white paper reframes data center efficiency around total energy consumed per unit of useful work rather than relying solely on Power Usage Effectiveness (PUE). PUE remains useful for describing facility overhead, but it does not directly indicate whether a facility completes more or less work per kilowatt-hour.

The operating premise is that, in modern air-cooled environments, thermal and mechanical behavior inside IT hardware can change the energy required to complete the same workload. Elevated inlet temperature, constrained thermal headroom, aggressive fan response, or vibration-sensitive storage behavior can extend runtime and increase total energy per completed job even when facility overhead appears improved.

The proposed optimized power management (OPM) framework uses server telemetry, workload telemetry, and facility telemetry to identify the operating region that minimizes energy per completed workload while maintaining ASHRAE environmental compliance and reliability limits. The concept is most relevant to air-cooled sites with compute-intensive or I/O-sensitive workloads, including mixed CPU/GPU clusters and HDD-based storage estates.

Because quantitative savings ranges in the source drafts are directional rather than independently validated, this consolidated draft treats expected benefit as a site-specific hypothesis to be measured through controlled pilot testing. The recommended validation method compares an intervention population against a control population using matched or normalized workloads, energy per completed task, throughput per megawatt, and reliability outcomes.

## Abstract

Power Usage Effectiveness remains useful for characterizing facility overhead, but it does not directly indicate how much total energy is required to complete a defined workload. In air-cooled data centers, thermal throttling, temperature-dependent leakage power, server fan-power escalation, and vibration-sensitive storage behavior can extend runtime and raise total energy per completed job. This paper presents OPM as a measurement and control framework that uses closed-loop telemetry from IT equipment, workloads, and cooling infrastructure to minimize energy per unit of useful work while respecting reliability limits and ASHRAE environmental guidance.

**Keywords:** air-cooled data centers; energy productivity; workload energy; PUE; thermal throttling; fan power; HDD vibration; supervisory control; ASHRAE TC 9.9

## 1. Introduction

For two decades, PUE has been the dominant facility-level indicator for data center efficiency. It remains useful for quantifying overhead associated with cooling and power distribution, especially when comparing similar facilities or similar operating conditions.

However, PUE does not include useful work output. It therefore cannot distinguish between two operating states that draw similar instantaneous power but complete materially different amounts of work over time. In modern air-cooled environments, that limitation matters because runtime can change as a function of inlet temperature, thermal headroom, workload type, fan behavior, and storage sensitivity.

The central thesis of this paper is that the relevant optimization target is not minimum facility overhead in isolation, but minimum total energy consumed per completed workload. This framing connects the physical data center to the business, scientific, or operational output being produced by the IT estate.

## 2. Problem Statement and Loss Mechanisms

Static temperature setpoints, airflow balancing, and PUE tracking remain necessary operating practices, but they are not sufficient to characterize energy productivity in compute-intensive or I/O-intensive IT assets. The consolidated source material identifies four interacting mechanisms that can increase total energy per job even when an infrastructure metric appears favorable:

- Semiconductor leakage current can rise with temperature and increase device power draw.
- Dynamic voltage and frequency scaling (DVFS) or thermal throttling can reduce throughput as hardware approaches thermal limits.
- Server fan power can rise nonlinearly with fan speed, making high fan excursions disproportionately expensive.
- Vibration-sensitive storage behavior, particularly in HDD-rich environments, can increase latency, retries, and workload duration.

Because these mechanisms affect both power and elapsed runtime, the system objective should be defined over an interval of completed work rather than from an instantaneous power ratio alone.

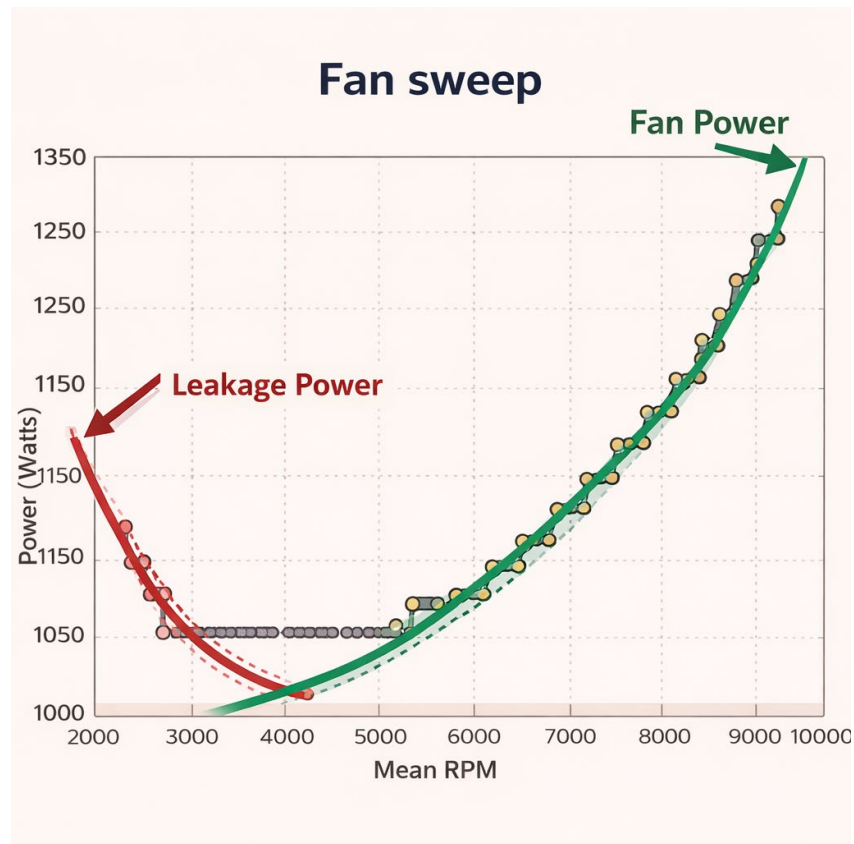


Figure 1. Fan-sweep graph illustrating nonlinear fan-power escalation and leakage-power sensitivity. The graph is illustrative and should be interpreted within the measured context of the source test conditions.

### 3. Why PUE Alone Is Insufficient

PUE is a ratio of powers. Energy consumption, by contrast, is the time integral of power across the duration required to finish a task. A control strategy that lowers cooling power but extends runtime through throttling, higher internal fan speed, or degraded storage performance may reduce PUE while increasing energy per completed job.

This distinction becomes more pronounced when server platforms are operated near the edge of their thermal or mechanical comfort zones. Leakage-aware management, fan-speed-aware scheduling, DVFS-performance interactions, and thermo-mechanical storage degradation all support the engineering premise that the relationship between ambient condition and total workload energy is nonlinear rather than monotonic.

Accordingly, PUE should be treated as a necessary but incomplete metric. For operational decision-making, it should be supplemented with an energy-productivity measure that couples facility behavior to completed work.

### 4. Proposed Metric Framework

The source drafts use phrases such as "work per watt" and "computational output per total facility power." For technical use, the metric should be defined in energy terms over a completed workload interval:

$$\text{Energy Productivity} = \text{Useful Work Output} / \text{Total Facility Energy Consumed}$$

Useful work output must be defined by the application domain. Suitable examples include completed simulations, training epochs, inference tokens, settled transactions, queries served, or terabytes processed. Total facility energy

should include, at minimum, the IT load plus the cooling and support energy attributable to the same test boundary and interval.

For comparison across conditions, workloads should be identical, replayed, or normalized to stable workload classes. Companion views such as throughput per megawatt, energy per completed job, throttling frequency, fan-power contribution, and storage-latency impact help determine whether a change improves energy productivity or merely redistributes power between IT and facility systems.

Table 1. Example useful-work units for pilot measurement

Environment	Representative useful-work unit	Primary outcome metric
AI training	Tokens processed or epochs completed	kWh per training unit; tokens per MWh
AI inference	Queries or tokens served	Wh per 1,000 responses; queries per MWh
HPC simulation	Completed jobs, solver iterations, or solved timesteps	kWh per job
Financial compute	Transactions settled or risk batches completed	Wh per transaction or batch
Analytics / storage-heavy	Terabytes scanned or completed batch runs	kWh per TB; latency at fixed throughput

### 5. Conceptual Energy-Optimal Operating Point

The control hypothesis behind OPM is that many air-cooled compute and storage populations have a workload-dependent operating region where total energy per completed workload is minimized. At colder conditions, overcooling can increase facility overhead. At warmer conditions, thermal stress can increase IT-side energy through throttling, higher fan speed, leakage-power growth, and longer runtime.

The optimum, therefore, need not coincide with the minimum PUE point. It should be discovered empirically for each site, hardware population, and workload class, then bounded by reliability requirements and applicable ASHRAE environmental guidance.

### 6. OPM Control Concept

OPM is framed as a supervisory, closed-loop layer. It is not a replacement for server firmware or building-management control. Instead, it coordinates available telemetry and selected control inputs to maintain a thermal-headroom band that minimizes workload energy while respecting reliability and environmental constraints.

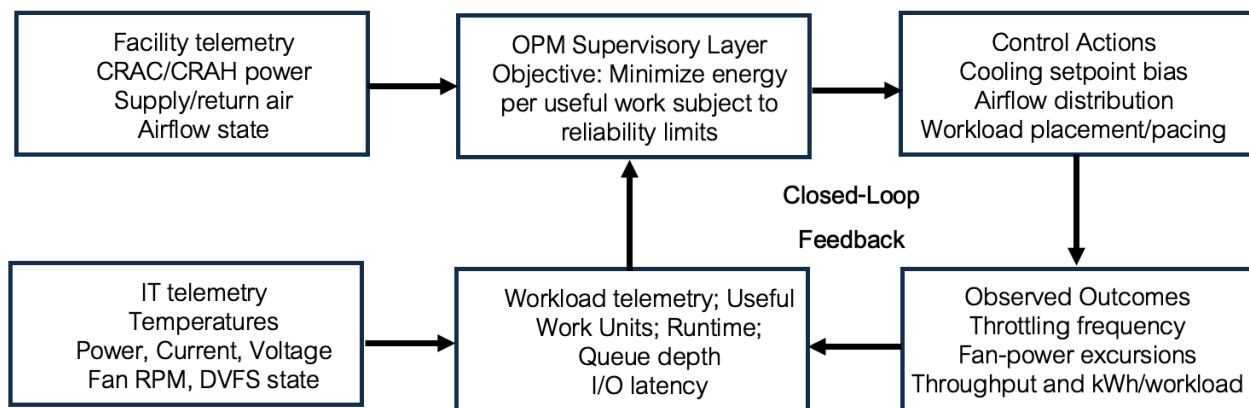


Figure 2. Supervisory control concept retained from the source material, linking facility telemetry, IT telemetry, workload telemetry, control actions, and observed outcomes in a closed loop.

Representative observations include server inlet and component temperatures, power and current telemetry, fan RPM, DVFS or throttling state, job runtime, useful work units, queue depth, storage latency, CRAC/CRAH operating

state, and cooling-system power. Optional mechanical indicators may include vibration or acoustic signatures associated with storage or fan-induced resonance.

Representative actions may include small biases to supply-air setpoint, airflow distribution adjustments, workload placement, and workload pacing when orchestration controls exist. From an engineering perspective, the significance of the concept is that it links facility control and IT performance in one measurement boundary rather than optimizing them independently.

## 7. Applicability and Boundary Conditions

The strongest applicability case is for air-cooled facilities with compute-intensive or I/O-intensive workloads, particularly where HDDs remain material to performance or where recurrent thermal throttling and fan excursions are already observable.

The concept is likely to be most valuable in the following conditions:

- Power-constrained air-cooled data halls where additional compute demand cannot easily be met with more infrastructure.
- Mixed CPU/GPU clusters with recurring thermal throttling or fan excursions.
- I/O-bound environments in which vibration or thermal effects measurably influence storage performance.
- Sites that already have reasonably efficient mechanical systems, where further gains depend on joint optimization of IT and facility behavior.

The expected benefit is likely to be smaller in fully liquid-cooled systems, highly tuned hyperscale data centers with mature workload orchestration, or facilities where the dominant inefficiencies lie outside the IT-thermal interaction described here. These boundary conditions should be stated explicitly so the framework is not over-generalized.

## 8. Measurement and Validation Protocol

For committee or owner acceptance, quantitative claims should be established by controlled measurement rather than point anecdotes. A practical validation sequence is:

1. Baseline period. Measure at least 30 days of steady operation for the target hall or cluster, including IT power, cooling power, inlet temperatures, component temperatures, fan speed distribution, throttling incidence, workload output, and reliability events.
2. Controlled intervention. Apply the OPM supervisory algorithm to a defined hall, pod, or cluster while maintaining a control population under existing operating practice.
3. Matched workloads. Use repeated or normalized workload classes so that differences in throughput and energy are attributable to operating conditions rather than workload mix.
4. Outcome metrics. Report energy per completed workload, throughput per megawatt, throttling frequency, fan power contribution, storage-latency impact, and any effect on service levels or reliability.
5. Acceptance criteria. Adopt only if the intervention improves energy productivity without causing adverse thermal excursions, increased error rates, or unacceptable equipment stress.

*Table 2. Recommended pilot evidence package for committee or owner review*

Phase	Primary measurements	Acceptance intent
Baseline	IT kW, cooling kW, inlet temperatures, fan RPM, throttling state, throughput	Establish representative operating range
Controlled rollout	Same measures on intervention and control populations	Isolate intervention effect
Matched workload analysis	Energy per completed workload, throughput per MW, storage latency	Normalize for workload mix
Reliability review	Thermal alarms, error rates, hardware events, SLA impact	Reject gains that degrade reliability

## 9. Treatment of Quantitative Claims

The source package contains several ranges for expected savings, including values described as directional. Such numbers may be achievable in specific high-friction environments, but they should not be presented to a standards committee as generally demonstrated performance unless accompanied by test conditions, sample size, instrumentation boundaries, workload-normalization method, and statistical treatment.

The more defensible formulation is that telemetry-driven optimization may materially improve energy per completed workload in suitable air-cooled environments, with site-specific benefit determined by baseline operating practice, workload sensitivity, storage architecture, and existing cooling efficiency.

## 10. Relationship to ASHRAE TC 9.9

ASHRAE TC 9.9 guidance emphasizes environmental envelopes, reliability, airflow management, and practical thermal design for mission-critical facilities. The OPM concept is best understood as a control and measurement methodology layered onto that guidance, not as a proposal to disregard recommended practice or to define a new allowable environmental class.

From a committee perspective, the contribution of this framework is methodological: it offers a way to combine facility and IT telemetry so operators can identify the most energy-productive operating point for a defined workload while remaining inside reliability and environmental constraints.

That distinction is important: the proposal is an operating methodology layered onto existing thermal guidance, not a replacement for environmental compliance.

## 11. Conclusions

The consolidated source documentation presents a coherent engineering argument that data-center efficiency should not be evaluated only by the ratio of facility overhead to IT power. In air-cooled environments with thermally sensitive compute or storage behavior, the more complete objective is minimum total energy per completed workload.

The OPM framework translates that argument into an engineering workflow based on telemetry, controlled intervention, and workload-normalized measurement. Its value lies less in any single savings claim than in the proposition that the energy-optimal operating point is measurable, workload dependent, and amenable to closed-loop control.

For ASHRAE TC 9.9 audiences, the next step is to standardize the measurement boundary, define useful-work metrics by workload class, and validate the control approach through instrumented pilot deployments.

## References

- [1] M. Zapater, O. Tuncer, J. L. Ayala, J. M. Moya, K. Vaidyanathan, K. C. Gross, and A. K. Coskun, "Leakage-power-Aware Workload and Cooling Management for Improving Server Energy Efficiency," IEEE Transactions on Parallel and Distributed Systems, 2014.
- [2] M. Zapater, J. L. Ayala, J. M. Moya, K. Vaidyanathan, K. C. Gross, and A. K. Coskun, "Leakage-power-Aware and Temperature-Aware Server Control for Improving Energy Efficiency in Data Centers," DATE, 2013.

- [3] S. Sondur, K. C. Gross, and K. Kant, "Thermo-Mechanical Coupling Induced Performance Degradation in Storage Systems," IEEE/ACM CCGrid, 2020.
- [4] J. Turner, "Effects of Data Center Vibration on Compute System Performance," USENIX Workshop on Sustainable Information Technology, 2010.
- [5] C. S. Chan, Y. Jin, Y. K. Wu, K. C. Gross, K. Vaidyanathan, and T. S. Rosing, "Fan-Speed-Aware Scheduling of Data-Intensive Workloads for Optimal Performance and Energy Efficiency of Data Center Assets," ISLPED, 2012.
- [6] K. Vaidyanathan and K. C. Gross, "Intelligent Power Monitoring and Management for Enterprise Servers," Berkeley IEEE Symposium on Energy Efficient Electronic Systems, 2013.
- [7] Rubenstein and Faist, "Data Center Cold Aisle Set Point Optimization," IEEE IThERM, 2014.
- [8] K. Vaidyanathan, K. C. Gross, and S. Sondur, "Ambient Temperature Optimization for Enterprise Servers: Key to Large-Scale Energy Savings," Berkeley IEEE Symposium on Energy Efficient Electronic Systems, 2015.

## Author Note

This manuscript is intentionally framed as a technical discussion paper based on the provided draft documentation. Quantitative performance claims should be treated as hypotheses to be validated through controlled site testing rather than as universally established outcomes.