TRUE NORTH PROGNOSTICS

EXECUTIVE OVERVIEW

Rethink Cooling to Boost Efficiency

Over the years, many data centers tried to cut costs by raising ambient temperatures to reduce cooling expenses. While lowering cooling can save some energy from reduced HVAC cooling costs, there's a much smarter way to achieve **much greater savings without hurting compute and I/O performance**.

Using **Optimized Power Management** (OPM) from TNP, energy use is continuously optimized across the entire data center. Instead of relying on one-size-fits-all settings, the system automatically finds the balance point where energy waste is minimized. [From excessive server internal fan speeds, leakage current in the CPUs/GPUs, and vibration-induced degraded I/O performance in spinning disk drives].

Rethink Cooling to Boost Efficiency

The result:

- **Lower overall energy costs**
- **Substantially boosts server performance (both compute & I/O performance)**
- **No additional hardware required**
- **Faster and more efficient operations**

For all modern enterprise servers, saving cooling energy by raising ambient temperatures in the data centers leads to CPU and GPU throttling, which results in substantially greater overall energy costs, while slowing down workload completion times.

TRUE NORTH PROGNOSTICS

Rethink Cooling to Boost Efficiency

Problem:

- Industry advice is to "warm up" data centers, resulting in hotter environments which increases server **fan speeds and vibrations**
  - All CPUs and GPUs today have frequency scaling (throttling). The degree of throttling goes up exponentially with temperature → **Increased job completion times for compute-intensive workloads**
  - For servers with spinning disk drives, increased fan speeds result in increased vibrations and much lower I/O performance → **Increased job completion times for I/O-intensive workloads**
  - For many enterprise servers, the fan motors now consume more power than the CPUs [the power for the fan motors goes up with the *cubic* power of the fan RPMs]

Rethink Cooling to Boost Efficiency

Solution:

Cooler temperatures make hardware run faster and more **efficiently**

OPM uses real-time telemetry data (e.g., distributed internal currents, voltages, temperatures, fan speeds) to:

- Calculate thermal headroom margins
- Dynamically set ambient temperature at an optimal (cooler) value
- Reduces fan speeds, minimizes chip throttling, significantly increases both CPU-intensive and I/O-intensive performance; reduces overall energy consumption across the data center

Rethink Cooling to Boost Efficiency

Savings:

- Up to 30% (or more) reduction of total electricity costs → (**O&M savings**)
- Faster job completion → **Improved equipment utilization**
- Fewer new servers and additional hardware needed → (**Capital savings**)

Appendix A

Warm Data Centers Waste Energy

Low PUEs save energy in the cooling assets, but by increasing ambient temperatures at the inlet of air-cooled servers, wastes far more energy in all the IT and Storage assets.

OPTIMIZED POWER MANAGEMENT ™

TRUE NORTH PROGNOSTICS

THE PROBLEM

- 2008 ASHRAE™ warm-room guidelines suggest decreasing HVAC energy costs^
- Hotter servers begin to throttle CPUs/GPUs → often doubling or more workload runtimes
- Aggregate fan-motor power rises cubically with fan speeds → steep energy penalties
- Increased fan speeds create more vibrations that slow I/O throughput → substantially increasing workload runtimes and wasting energy
- **Net result**: More Megawatt-Hours are consumed across the data center, not less

^ https://www.energystar.gov/products/data_center_equipment/5-simple-ways-avoid-energy-waste-your-data-center/raise-temperature?utm

OPTIMIZED POWER MANAGEMENT ™

OPM™ OPTIMIZES BOTH PERFORMANCE AND ENERGY EFFICIENCY

Key Takeaway:

Cooling data centers intelligently with OPM™ saves far more energy than warm-room approaches.

By preventing throttling, Improving I/O, and slashing fan power, OPM™ optimizes both performance and efficiency.

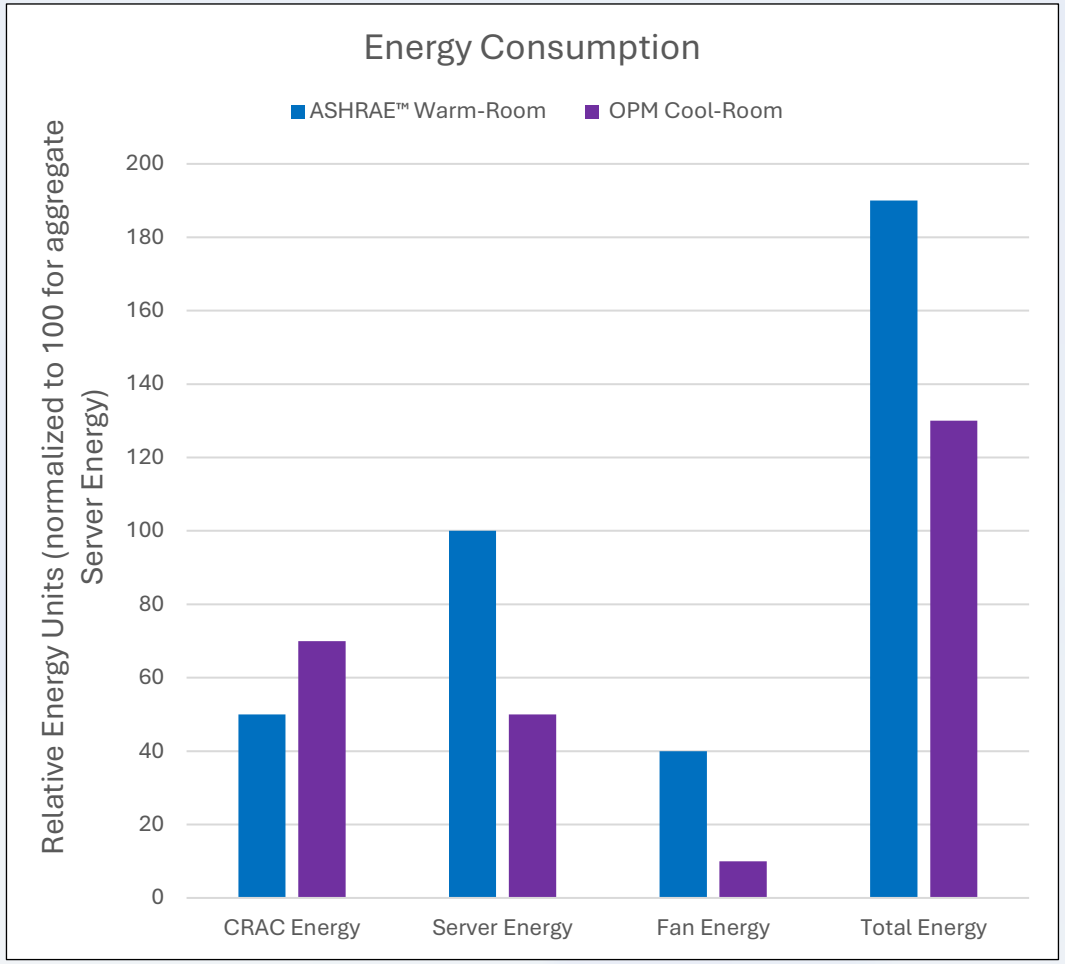

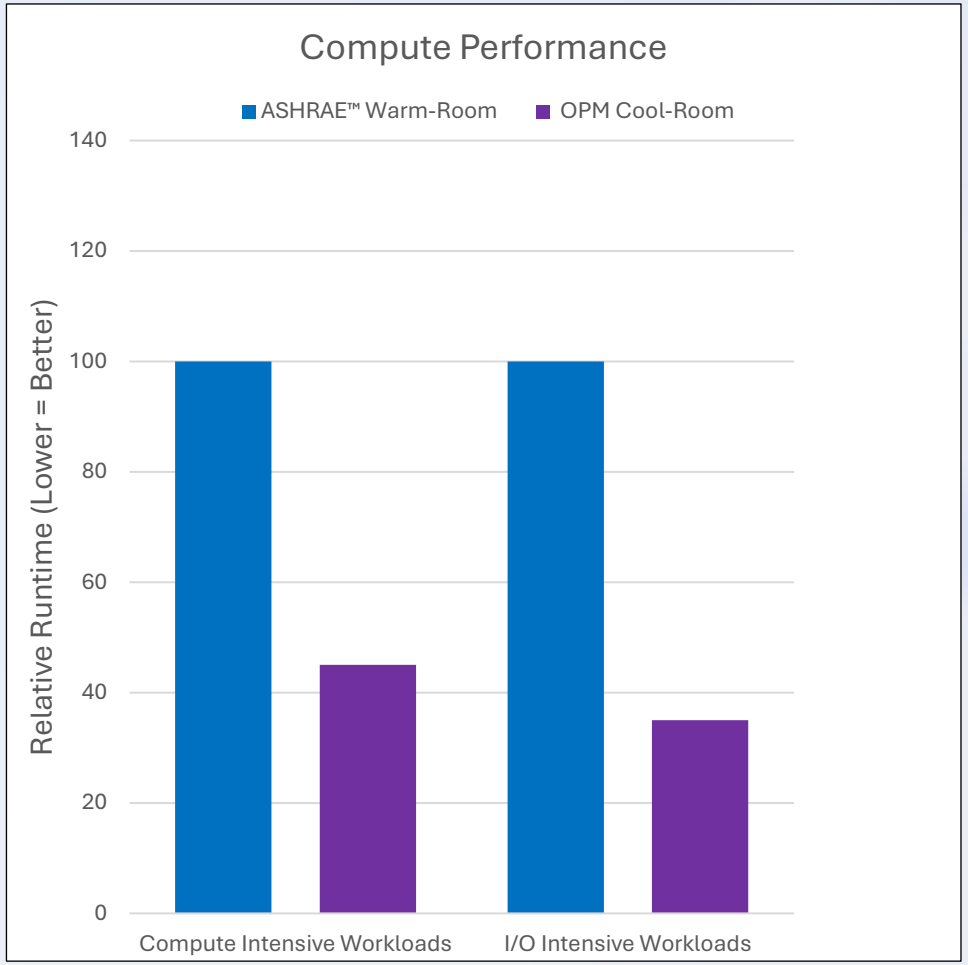

ASHRAE™ - American Society of Heating, Refrigerating and Air-Conditioning Engineers

TRUE NORTH PROGNOSTICS

HISTORY OF PUE™ – POWER USAGE EFFECTIVENESS

$$\text{PUE}™ = \frac{\text{Total Facility Energy}}{\text{IT Equipment Energy}}$$

In 2006, the Green Grid Consortium (Intel, AMD, HP, Microsoft, etc.) introduced Power Usage Effectiveness (PUE). Defined as the ratio of Total Facility Energy to IT Equipment Energy, the ideal value is 1.0

Hyperscale operators like Google and Microsoft began publishing PUE numbers. The industry average was 2.0–2.5, while Google reported ~1.2 {DeepMind}. PUE became a benchmark in sustainability reports and a marketing tool, and standards bodies such as ASHRAE and the EU Code of Conduct for Data Centres began adopting PUE.

Hyperscale leaders have achieved fleet-wide PUE averages around 1.10–1.12.

HOWEVER: PUE does not measure the data center work (compute transactions and cumulative I/O) being performed, i.e. “Work per Watt”.

By metering the cumulative compute-and-IO work being done, it is discovered (and proved quantitatively for any modern air-cooled servers with any CPUs and/or GPUs), that warming up the data center (to lower the PUE) *wastes more energy in the IT systems than is saved in the cooling systems.*

HISTORY OF PUE™ AND WHY IT'S NOT A MEANINGFUL MEASURE TODAY

Reducing PUE™ was universally a good objective with computers and storage from 20 years ago and earlier. Before the mid 2000s, all computer chips ran at the same operating frequency regardless of ambient temperature and I/O rates for spinning disk drives were unaffected by fan vibrations inside the servers and storage arrays. Moreover, almost all server fans were constant speed fans, and the fan motor energy was a small portion of the energy budget for servers and storage.

TRUE NORTH PROGNOSTICS

HISTORY OF PUE™ AND WHY IT'S NOT A MEANINGFUL MEASURE TODAY

Today, with modern computers and storage, reducing PUE™ by raising ambient temperatures is extremely counter productive because:

- All CPUs and GPUs have dynamic frequency scaling, and the "throttling" goes up exponentially with chip temperatures
- All spinning disk drives have become hyper-sensitive to low-levels of vibrations. As ambient temperatures warm up in the data center, server and storage fans increase their RPMs, raising rack vibration levels, substantially reducing I/O rates
- Thanks to Moore's Law, fan motors (and the numbers of fans) have gone up, and for all fans, the energy for the fan motors goes up with the cubic power of the fan RPMs. Raising the data center ambient temperature significantly increases the "energy wastage" from the fan motors ... and all of the heat dissipated from the fan motors incurs a "double penalty" (it adds to the heat that must be removed by the CRACs)
- Most significantly (and not reflected in the PUE™ metrics): when the compute performance for the IT systems is cut in half [from frequency throttling in the chips and I/O degradation in the disk drives], it means that for fixed customer workloads, the completion times double, burning 2x the energy for everything in the servers [memory, power supply units, PCIE I/O cards, ASICs, fan motors, as well as the CPUs and HDDs]

History of PUE and why it's not a meaningful measure today. . .

While reducing HVAC cooling costs does improve the PUE [as was a great goal from the 1960s-1990s], for todays enterprise IT systems, warming the data center wastes far more energy [and generates far more carbon costs], than cooling the data centers back down.

The present ASHRAE guidelines for computing PUEs does not reflect the fact that the "work done" by the IT systems [the cumulative compute transactions and I/O rates] are substantially degraded by warming the data center. By cutting cooling costs to reduce PUEs, the total energy costs to complete customer workloads is substantially increased [and the data center's "customers" running the workloads have to wait significantly longer for workload completion times]. Moreover, the organizations have to buy significantly more hardware for the data center to meet business "capacity planning" criteria.

By cooling the data center, overall data center energy is reduced, all compute and I/O performance for the IT assets are significantly boosted.

With no hardware modifications anywhere in the data center, OPM's AI_MSET identifies the optimum ambient temperature that substantially boosts both compute and I/O performance, shortens workload completion times, and significantly reduces the total energy requirements across the data center.

TRUE NORTH PROGNOSTICS

Conclusion:

ASHRAE thermal guidelines advocating warmer data centers to save energy are well meaning, but now-days are very mis-guided and counter productive (except for data centers in year-round cold climates where it becomes possible to leverage free air cooling.

Doing so will result in very low PUE ratios for the data centers, but the fallacy in associating PUE ratios with energy efficiency lies in the fact that for all recent server and spinning-storage technology, warming up the data center will severely degrade both CPU and I/O performance, while degrading IT energy efficiency by 50%-80% in terms of “work done per energy consumed, (i.e., Work per Watt™).

By cooling the data center to a value that minimizes the insidious penalizing energy-wastage mechanisms that are present in all air-cooled enterprise server and spinning-disk storage systems today, substantial energy is saved, significantly increasing overall compute performance, decreasing the carbon footprint for the data center owner/operator, while increasing return-on assets for the IT assets throughout the data center.

TRUE NORTH PROGNOSTICS

Optimized Power Management: Rethink Cooling to Boost Overall Efficiency

Work per Watt™

- OPM saves tons of carbon, but at the same time also saves Megawatts (lots of money). It also significantly boosts both compute performance and I/O performance resulting in substantially greater ROI for the capital investments in the data center equipment
- Today, the PUE measurement has become counter-productive when chips started throttling and when storage became hyper-sensitive to vibration levels over the past 15 years

Optimized Power Management™ Bibliography

External Scientific Publications documenting the physics behind the three "Parasitic Energy-Wastage Mechanisms" that are triggered by warming the ambient air for modern IT air-cooled server and storage assets: CPU/GPU Leakage Power, Fan Motor Power, and System Ambient Vibrations.

- "Simware: A Holistic Warehouse-Scale Computer Simulator," Georgia Tech Journal Paper, IEEE Compute Journal (0018-9162/12), 2012. [Shows that as data center temperature rises, fan motor power in the IT assets rises significantly, wasting energy and distorting the facility PUE.]
- “Thermo-Mechanical Coupling Induced Performance Degradation in Storage Systems,” S. Sondur, K. C. Gross, and K. Kant, 20th IEEE/ACM International Symposium on Cluster, Cloud, and Internet Computing, Melbourne, Australia (May 11-14, 2020).
- “Leakage-Aware Workload and Cooling Management for Improving Server Energy Efficiency,” M. Zapater, O. Tuncer, J. L. Ayala, J. M. Moya, K. Vaidyanathan, K. C. Gross, and A. K. Coskun, Transactions on Parallel and Distributed Systems Journal (June 29, 2014).

Optimized Power Management™ Bibliography, cont'd

- "Effects of Data Center Vibration on Compute System Performance," Julius Turner, Proceedings of USNIX Workshop on Sustainable Information Technology, San Josa, CA (2010) *[Shows 85% degradation in IO rates for IO-intensive workloads in on-prem data centers]*
- "Intelligent Power Monitoring and Management for Enterprise Servers," K. Vaidyanathan and K. C. Gross, 3rd Berkeley IEEE Symposium on Energy Efficient Electronic Systems (E3S-2013), Oct 28-29, 2013.
- "Impact of Dynamic Voltage and Frequency Scaling on n-Tier Application Performance," GA Tech and Fujitsu Laboratories, Proc. of ACM Conf. on Timely Results in Operating Systems (TRIOS'13), Nov 2013.
- "Leakage-Aware and Temperature-Aware Server Control for Improving Energy Efficiency in Data Centers," M. Zapater, J. L. Ayala, J. M. Moya, K. Vaidyanathan, K. C. Gross and A. K. Coskun, International Symposium on Design Automation and Test in Europe (DATE-2013), Grenoble, Fr. (Apr 2013).
- "Ambient Temperature Optimization for Enterprise Servers: Key to Large-Scale Energy Savings," K. Vaidyanathan, K. C. Gross, S. Sondur, 4th Berkeley IEEE Symposium on Energy Efficient Electronic Systems (S3E-2015), Berkeley, CA (Oct 28-29, 2015).

Optimized Power Management™ Bibliography, cont'd

- "Data Center Ambient Temperature Optimization for Large-Scale Energy Savings," K. C. Gross, Invited Paper and Keynote Industrial Presentation, EcoCloud Annual Symposium on Enterprise Computing Energy Optimization, EPFL, Switzerland (June 2015).
- "Fan-Speed-Aware Scheduling of Data-Intensive Workloads For Optimal Performance and Energy Efficiency of Data Center Assets," C. S. Chan, Y. Jin, Y. K. Wu, K. C. Gross, K. Vaidyanathan, and T. S. Rosing, Proc. Int'l. Symp. on Low Power Electronics and Design (ISLPED12), Redondo Beach, CA (Jul 2012)
- "Correcting Vibration-induced Performance Degradation in Enterprise Servers," C. S. Chan, B. Pan, K. C. Gross, K. Vaidyanathan, and T. S. Rosing, Proc. ACM 2013 International Green Metrics Conference, Pittsburgh, PA (Jun 2013)
- "Data Center Cold Aisle Set Point Optimization," Rubenstein and Faist, Microsoft Research, Proc. IEEE ITHERM Conference (June 2014).