

# INCREASING DATA CENTER ENERGY EFFICIENCY BY MONITORING AND TARGETING TO HEATING AND COOLING METRICS

*A White Paper*

*By: Robert Hunter and Chet Sandberg*



## Abstract

Data center heat density has been increasing since the advent of the server. This has become particularly problematic during the past few years as data center managers have struggled to cope with heat and power intensity problems. These issues have resulted in enormous energy bills and a rising carbon impact. The purpose of this paper is to examine the root cause of this problem and then determine a logical course of action to minimize energy use and consequent environmental effects of data center operations while maximizing processing throughput and system uptime.

We will see that trying to use fixed cooling assets to try to address a dynamic heat load creates inherent inefficiencies. Standardized metrics of wattage are detailed for all IT and cooling loads. With this data, dynamic loads can be converted from the source of the problem to a solution by using new tools to balance intra-cabinet loads and lower total energy costs. This design will be shown to produce the highest return on investment for increasing the energy efficiency in a data center.


## Background

### Moore's Law

Computer servers are empowered by CPU chips whose advance in capabilities are generally described through Moore's Law. Moore's Law states that the number of transistors on a chip will continually double every 2 years. This law has been much debated over the past few years but, as the graph in Figure 2 shows, it has proved remarkably prescient.

The fact that Moore's Law continues is all the more remarkable because of the enormous power challenges that exist in the design of chips. This power challenge is posed by what we describe as the Corollary to Moore's Law which is: "The general trend is for maximum power consumption to increase a little more than 2 times every 4 years."<sup>1</sup> Much of the engineering resources used by chip manufacturers today are spent tackling this power consumption and a related heat dissipation challenge.

Figure 1



Before proceeding, it is important to define the units of measurement of CPU power, heat and total energy used. All of these items are integrally related. The power that a CPU draws, in watts, is exactly the same as the heat it radiates in watts. That is, power draw and heat load are simply two sides of the same coin. Wattage is defined as the instantaneous measurement of heat within a circuit whereas, what is commonly referred to as energy, is the accumulation of wattage data over any time period.

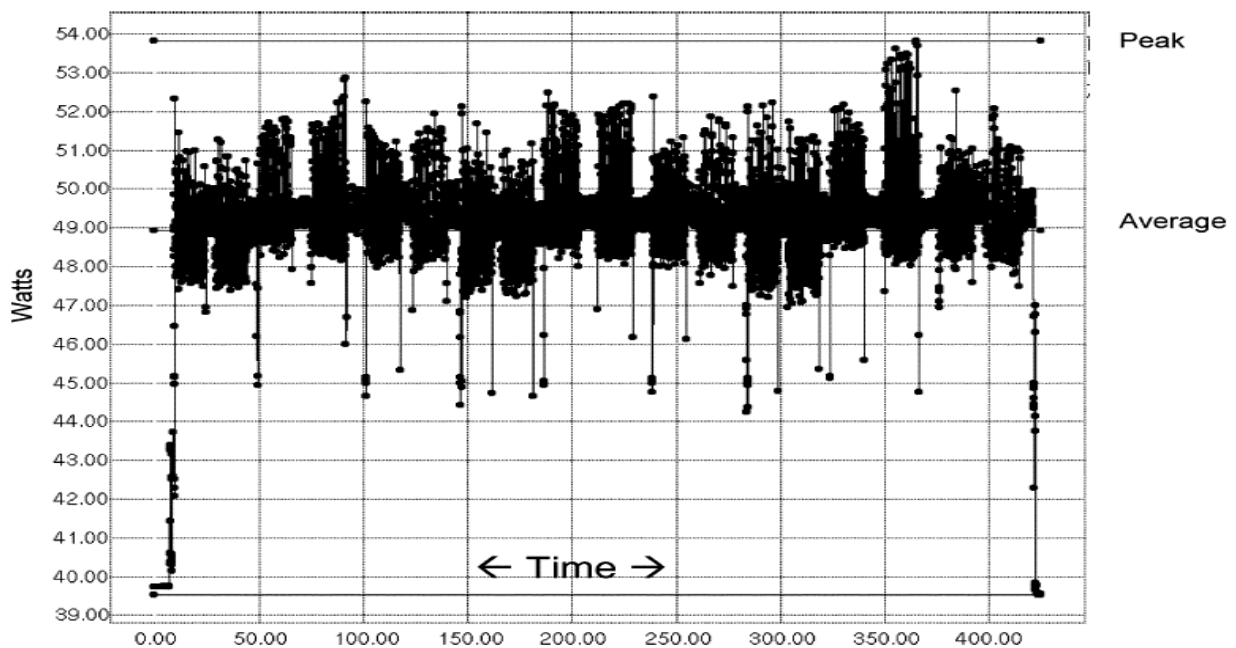
### Intra and Inter-Chip variability

As the number of transistors increase along the path predicted in Moore's law, the nano scale of circuits continues to shrink for CPU chips. This nano scale brings in many challenges, most notably in terms of variations within cross sections of circuitry among otherwise identical chips.

Intra-chip variability has been highly scrutinized, especially with the advent of multi-core chips. In a multi-core architecture, each core forms a separate processing zone. Humenay et al examined the variability of wattage and other parameters within a dual core chip related to circuitry variations.<sup>2</sup> His team concluded that asymmetry between any two cores was systemic and that asymmetry leads to differences in wattage drawn by a chip to perform the same task.

Figure 2 depicts the effect of variability on wattage within the first-generation dual-core Itanium processor (Montecito) running a single task repetitively.<sup>3</sup> As this graph shows, the power consumption in watts has a range of around +/- 8 % from the mean while running the same task. This illustrates the fact that a single chip will use differing amounts of energy to accomplish the same process.

**Figure 2**



A second area of chip variability occurs between any two chips made with the same die. This is known as With-in Die (WID) variability. WID variability, has been studied with great detail over the past few years as nano-scale architectures have moved to new levels. Studies have shown that wattage variability for chips produced on one die can have standard deviations of +/- 10% from the mean.<sup>4,5</sup> That is, 68% of the chips in the same die would have an *average* wattage draw that falls within a range of 20% from top to bottom. For the second sigma group of 27%, one can

expect a power range of 40% and so on. Such a high range of power consumption is far beyond what many have come to expect and it creates new management challenges.

There is a third group of chip variability which is known as the Inter-die variation. Today, this area has been largely unexplored with detailed public research. Information concerning inter-die variations remains a closely watched number for chip manufacturers as it directly affects their yield ratios of chip production. While we will not speculate on specific numbers in this paper, one can say that two chips from different dies would appear to have a likelihood of producing a greater variation than those that come from a common die.

In addition to the variation between any two identical chips, one must deal with the natural change in task-driven wattage loads during the day. These variations typically go from an idle position for at least some time of the day towards a full load on some servers. The variation of power usage between identical chips coupled with the natural variation of work load provides an extremely dynamic heat profile for each and every server. This contributes to system cooling inefficiencies, which need to be dealt with if energy use and carbon impact are to be reduced. .

Two key issues face data center managers today:

- Managing dynamic heat within data center racks through the use of fixed cooling assets
- Managing rising energy costs and related carbon impacts related to these loads

The rise in magnitude and variability of heat loads creates enormous strains on data center cooling infrastructure. Heat problems manifest themselves as hot spots, hot zones, tripped breakers and other density-related issues. Rising energy costs are related to both the rise in magnitude and the inefficiencies created by load variability. This paper will examine the management of cabinet heat loads together with the matching of cooling assets to cabinet heat loads with the goal of increasing total energy efficiency

### **Establishing Data Center Management Metrics from RMS Wattage**

The principal of Management By Information (MBI) states: "You can only manage what you measure." This principal is as vital to controlling data center efficiency as it is to manufacturing and process efficiency. In the past few years, The Green Grid and other organizations have proposed using standardized measurement metrics for overall data center energy efficiency.<sup>6</sup> The specifics of the metrics offered by the Green Grid are centered around wattage data and include:

- Power Usage Effectiveness (PUE):  $\text{Total Data Center Power} / \text{IT Equipment Power}$ .
- Data Center Efficiency (DCE) the inverse of this metric or;  $\text{IT Equipment Power} / \text{Total Data Center Power}$

Total Data Center Power includes the following:

- IT equipment – servers, storage, network equipment, etc
- Cooling Load – Chillers, CRAC (Computer Room Air Conditioners)
- Electrical losses associated with the PDU (Power Distribution Units), UPS (Uninterruptible Power Supplies) and Switchgear systems.

According to the Green Grid, most data centers have a PUE of over 3, yet a number of less than 1.6 has been shown to be achievable. Intel Corporation showed that, using an advanced strategy employing wet-side economizers, a number of as low as 1.3 is actually achievable.<sup>7</sup> In a PUE of 3, the cooling load exceeds the load of the IT load plus electrical losses. Thus, with a PUE of 3 a data center's cooling energy consumption can exceed the IT equipment energy usage.

By applying the principal of Management By Information to the challenge of data center efficiency management, it is of critical importance to measure true RMS wattage for all IT equipment, electrical equipment and cooling loads. There are several choices for measuring loads at the circuit level. These include:

- Amperage-only monitoring
- Derived or Estimated Wattage monitoring
- True RMS Wattage monitoring

The RMS stands for root mean squared. The RMS of a circuit's wattage is calculated by taking the square root of all of the positive and negative points on each voltage and amperage sine wave, averaging the results and providing the square of that average. True RMS data provides the actual heating equivalent value for any circuit. In order for the heating value of a circuit to be expressed in watts, it *must* incorporate the true RMS values for voltage and amperage and the power factor of the angle between the two. The result is a standardized measure that forms the basis for all measured electricity billing throughout the world which works for any sine wave of power, no matter how distorted it may be.

The heating value of true wattage is vital to understanding the actual heat produced by any server or piece of networking equipment. It is also vital to the sizing of cooling loads in a data center. The BTU (British Thermal Units) are defined as the amount of heat that it takes to change 1 pound of water by 1 degree F) required to be removed by the cooling system has a direct mathematic relationship to wattage. That relationship is expressed as:

$$1 \text{ kWh} = 3,413 \text{ BTU's of Heat}$$

The use of averaging methods to measure circuit power parameters uses the fact that RMS value is 10.99% more than an average for a *pure sine wave*. This less expensive form of measurement is sometimes employed as a low-cost monitoring alternative. However, *any waveform that is not a perfect sine wave will break the mathematical relationship between average and RMS and render an inaccurate result*. The more distorted a wave form, whether by harmonics or general interference, the more the error. Since all waveforms within a data center have some degree of distortion and since all data centers have some degree of harmonics, averaging sensors will inherently produce an incorrect answer vis-a-vis the true RMS wattage. Error rates can be very significant.

Some data centers employ first-generation monitoring devices that are included within power strip PDU's or rack PDU's. Most of these devices monitor average amperage values and either estimate voltage and power factors to estimate wattage values within a circuit. Clearly, using averages and/or estimates can provide readings that create a misleading picture for any data point. RMS readings must also be sampled *and* accumulated continuously at the sub cycle level in order to generate the actual heat balance and energy as measured in kiloWatt hours. The sheer volatility of CPU load makes this a necessity. Therefore, if one measures power and heat, it is imperative to have quality data RMS that is continuously sampled.

Gathering wattage data for a CRAC unit is not possible from a PDU as CRAC units are sufficiently large so as to have their own power breakers within a panel. The principals of Management By Information provide that one should never mix and match sources of information – in this case estimates of wattage from a PDU strip vs actual wattage and BTU requirements for CRAC units. We therefore propose that the most effective and accurate point of measurement for both cabinet loading and cooling system load is at the circuit-level within each power panel. We further

propose that each circuit's information can logically be assigned to its usage (servers within a cabinet and their users and CRAC and chiller units) via a relational database. Software tied to the relational database can read the real-time data from each true RMS IT and cooling load.

The use of a relational database provides a number of valuable features including:

- Wattage data by circuit can easily be combined to see total heat wattage by cabinet
- Cabinet heat loads to be matched against individual CRAC unit cooling resources
- What-if scenarios can be employed by moving circuits virtually within a floor space to see the effect on heat and cooling efficiencies before a hard move of devices is completed
- Energy Efficiency Ratio (EER) can be seen as trends as is discussed in detail later in this paper.

Because power circuits are fed conveniently from panels that contain large groups of circuits, it can be much more economical to measure circuit-level wattage in these neatly grouped units. Typically, power panels consist of 42, 48, 84, 98 or even 100+ circuits. The ability to measure large groups of circuits from a single unit creates significant economies of scale vis-a-vis measuring circuit within a cabinet one power strip at a time. Monitoring at the panel-level, also allows the accuracy of measurements to reach utility-grade levels while maintaining a cost that can be considerably lower than PDU strip monitoring. Highly accurate current transformers, voltage transformers and utility-grade power meter chips can be employed. This is especially advantageous as companies begin to charge-back for individual energy usage within data centers. All of this data can be available instantaneously to form a true high-low-mean profile on a sub-minute basis.

### **Measuring and Managing Heat at the Source, the Cabinet-level**

Server to server variations in power usage have been shown to be substantial. But, the power usage of any individual server will be the same regardless of where it is physically located within a rack. It is the placement position of that server within a rack that greatly effects the path of its heat circulation as well as the circulation patterns of nearby servers. This change in circulation patterns, in turn, creates enormous differences in the amount of energy required to cool that server and other servers within a rack.

A recent study by Patel et al on the energy effect of cabinet placement options reviewed the relative effectiveness of heat dissipation strategies with a data center. It concluded that the majority of all data center system exergy loss (a measure of system efficiency) is attributable to the air flow patterns within individual data cabinets and between groups of cabinets.<sup>8</sup> This may be counter-intuitive and is opposed to common heat and cooling management practices.

Adjusting the air flow under floor, above floor or within a cabinet is a primary practice to provide more cooling to cabinets that are running at high heat levels. Adding additional air flow can offer help in some cases but at a very high price of extra energy usage. Air flow must increase at an exponential rate in order to dissipate an arithmetic rise in heat load. This means that the energy consumption to generate higher fan flow rates increases at an exponential rate as well.

A significant opportunity to lower cooling costs lies with the heat loading profile within each individual data cabinet. The study of cabinet heat load balancing can be divided into two separate categories, inter-cabinet and intra-cabinet. In a study focusing on inter-cabinet load placement, Norota showed that the more uniform the distribution of heat load, as measured in watts, the higher the energy efficiency achieved to cool that load.<sup>9</sup> Patel, et al showed in another study that a continuously calibrated placement of heat loads across the data center floor can provide up to

17% in energy savings from cooling<sup>10</sup> Later work by Patel, et al showed that up to 33% savings could be achieved by inter-cabinet balancing if work load could be scheduled so that computers and CRAC units could be turned off for portion of the day. While such a possibility may not be viable for many data center managers, one can view their work projecting 17% savings as a target for inter-cabinet balancing.

One key assumption in the studies of inter-cabinet balancing was that all intra-cabinet loads had already been balanced. That is, these studies looked at each cabinet as having its own uniform heat profile. In practice, this is seldom done but, it makes the possibility worth investigating as to the practical value of intra-cabinet balancing if it can be achieved.

Rambo and Yoshi undertook a study on intra-cabinet heat balancing and demonstrated that proper balancing can reap significant savings.<sup>11</sup> This study looked at managing individual servers by measuring the wattage of each and locating them in the best possible position. While the results of this study were promising, the practicality of measuring each plug-load at the RMS level is not presently economically feasible. However, Rolander, working with Yoshi on a follow-up study, took this cabinet balancing concept one step further toward a practical solution. In his study, Rolander showed that significant cooling savings can be achieved by breaking up servers into just 3 individual heat zones within a cabinet. This work, of some 300 pages of detail, showed that a net energy savings of 33% could be achieved or, alternatively over 50% more power load can be reliably maintained at an appropriate temperature.<sup>12</sup>

Further, Rolander showed that proper arrangement of the 3 groups provides from 20% - 60% less intra-cabinet temperature variability, a huge point for variable hot spots within cabinets that have been so illusive to control due to chip load variability. In discussing the results of the study, Rolander demonstrated that without intra-cabinet load balancing, the use of increasing air flow and cooling is actually very ineffective as an efficient cooling strategy. In fact, increasing air flow can actually have a negative effect on controlling temperature. Thus, trying to adjust the fixed fan units as the first line of dynamic heat mitigation is ineffective as compared to dynamic heat and power balancing on an intra-cabinet.

The choice of 3 measured groupings by Rolander, corresponds well to the number of circuits within an individual data cabinet. While some data centers employ 2 circuits per cabinet, an increasing number of data center managers bring from 3 to 4 circuits to each cabinet. This creates a significant opportunity to use the circuits as natural zone groupings within each cabinet for the purpose of actively managing each grouping by RMS watts of heat.

These studies have shown that a very good opportunity exists to manage data center heat and lower cooling if one can view and manage server heat loads by groups, at the circuit level within a cabinet or rack. From these intra-cabinet studies, one can also conclude that, the dynamic nature of server loading can be used as a strategic asset for intra-cabinet loads. Further, by using circuit-level monitoring, one can view live and historical heat loads and trends for groups of servers within each cabinet as part of a MBI efficiency plan. The ability to balance groups of heat loads can lead to reductions in cooling costs and thereby lead to lowering overall data center energy costs as it increases the PUE value. One can see that intra-cabinet balancing can be achieved by grouping servers within each cabinet in to as few as 3 groups and, we would suggest, to as little as two for practical reasons.

Inter-cabinet balancing can also reap some rewards. However, such inter-cabinet balancing can require significant computing resources to calculate appropriate loading changes via Computational Fluid Dynamics (CFD). The ability to make changes on an inter-cabinet basis are

greatly challenged by the fact that large groups of diverse users have non-transferable ownership of individual cabinets. Thus, while such balancing is possible it may be much more difficult to achieve than inter-cabinet balancing, where one of several owners already control the ownership of that cabinet and have incentive to reduce its total cost of ownership (TCO).

### **Measuring and Managing Cooling at the Source: The CRAC and Chiller Units**

We have seen that intra-cabinet heat balancing can be employed to maximize cabinet airflow patterns and reduce cooling cost. However, the fact remains that there is only one place in a data center where heat is actually *exchanged*. This exchange takes place only within the CRAC and Chiller system. The ability to match the cooling output of individual CRAC units to the actual heat generated within each data cabinet is *the* vital cog in the MBI management plan for a data center.

CRAC units use large amounts of energy for small changes in temperature settings. Adjusting CRAC temperatures by just 1 degree downward costs an additional 4% in energy usage<sup>13</sup> Unfortunately, because one always tends to cool to their hottest cabinet, this is a very common practice in many data centers to attack hot spots by simply lowering the supply temperature of one or more CRAC units.

Just as cabinet heat balancing can lower net cooling costs, so too, matching cooling requirements to actual heat loads can produce large energy savings. A significant number of data centers have been shown to actually over-cool their data cabinets.<sup>14</sup> Over-cooling can be the result of several things, most notably initial designs that, quite correctly, call for building-in load growth and redundancy. But, in day-to-day operations, a user that has access to cabinet heat data can assign individual CRAC units as primary cooling sources to groups of cabinets. In a study to be published by the authors, a savings of 8% in cooling load was achieved at a large UK data center simply by matching CRAC cooling to cabinet heat as measured in wattage.

In addition to energy savings from matching heat and cooling, the management of CRAC unit supply and return air temperature balancing can provide significant energy savings. CRAC units operate more efficiently when the return air (the air arriving at the CRAC unit from the data center) is sufficiently higher than the supply air (the cooled air leaving the CRAC unit). This temperature difference between supply and return is known as  $\Delta T$ . In general, the higher the spread in  $\Delta T$ , the higher the efficiency and visa versa. Concentrated heat loads arriving at the return air side provide a higher  $\Delta T$  and therefore, higher energy efficiencies.

The concept of hot-isles and cold isles is one example of a concentrated heat load strategy.<sup>15</sup> While such strategies are promising, they must take care to avoid mixing hot air and cold air to be effective.<sup>16</sup> Other heat concentration strategies include using hooded exhausts ducts at the cabinets and using cabinet-mounted CRAC units. The more efficiently one is able to contain and move the exhaust heat from an IT cabinet to the CRAC unit, the more efficient the cooling process will be. Yet, the success of any strategy can only be judged by measuring the energy used in the cooling system before a project is initiated and then after the project is implemented on-going basis.

Having a baseline for CRAC energy usage, by unit, and having a continuous feed of live wattage data from each unit is critical to measuring and maintaining the success of any cooling strategy. The dynamic nature of heat loading implies that, even good strategies may eventually lose their effectiveness over time and new strategies will need to be employed. Throughout the process, the only metric that can be relied on is the energy efficiency of each CRAC unit and of the entire cooling system.

Ironically, while cooling systems use the majority of energy in most data centers, almost no cooling systems are monitored for energy usage. We propose an affordable and understandable means to change this. We specifically propose the use of the Energy Efficiency Ratio of cooling loads as a standard to measure each CRAC unit and chiller unit's efficiency. EER is a metric established to determine the efficiency of each piece of cooling equipment to remove heat and is commonly used for HVAC equipment in office buildings and even homes. The calculation of EER for any piece of equipment is as follows:

$$\text{BTU's of cooling capacity} / \text{Watt Hours used in cooling}$$

The EER quoted by manufacturers is a theoretical value at an outside temperature of 95 degrees. We suggest that it is also an excellent tool to compare between CRAC units using live and historical data. By comparing the actual BTU's of heat removed with the amount of energy in kWh used to remove this heat, a common metric can be established to measure performance of each CRAC unit over time. This same data can be used as well as to compare one chiller to another or to compare a chillers efficiency over time.

The absence of an existing metric to evaluate each CRAC unit and the chiller system serves as a significant obstacle to effectively manage the largest segment of energy usage in a data center. The establishment of such a metric provides the basis to begin the management process in a structured method.

We therefore propose that each CRAC unit and each chiller unit should be monitored for BTU's removed and kiloWatt hours used in the removal process. The ability to trend such data during each day as part of an energy reduction strategy is vital to providing the benchmark for success to any energy efficiency project.

We further propose that the wattage of cabinet heat be matched to each CRAC unit to increase cooling system efficiency. By trending the actual heat value in watts for each cabinet over one month's time, a reliable figure can be used to forecast heat usage going forward. With this correct heat value, the customer can employ adjustable vent floor tiles and other adjustable air support structures to allow each cabinet to receive the exact amount of cooling required. This process reduces over-cooling and can show very quick payback as an implementation strategy. In addition, by continuing to track cooling in any cabinet, alarms can be set on wattage to alert the user as to when additional cooling or reduced cooling may be used to further target cooling resources to heat loads.

Having wattage data for each cabinet to match heat with cooling and having EER data available for each CRAC unit and chiller forms the basis for a sound MBI process for any sized data center. While the energy savings reaped by large data centers in these processes are clearly significant, the actual percentage energy savings on smaller data centers can actually be bigger due to the inherent inefficiencies present in many small data rooms.

The ability to establish and trend true RMS heat and cooling data is a necessity to control energy usage. In the absence of proper data, no established metric can be effectively employed and no energy savings from a product or strategy can be properly judged as to its success. However, the continuous monitoring of kiloWatts and kWh of heat, together with kiloWatts, kWh and BTU's removed by cooling resources, allows specific and gaugeable metrics to be employed to judge the effectiveness of any energy and carbon reduction strategy that effects the IT and cooling loads within a data center.

## **Summary: A Step by Step Approach to Data Center Energy Efficiency**

In this paper, we have proposed a set of metrics and tools to manage heat and thereby increase energy efficiency within the data center. Metrics are shown to gauge the energy efficiency of any data center. Measuring wattage for IT and cooling loads by circuit with high accuracy is shown as imperative and affordable.

We have shown that server heat loads are both increasing and are becoming more variable. Of the two, the variability creates the most energy inefficiencies. Load variability can be mitigated, and thus the standard deviation be lowered, by increasing the loading of individual servers through virtualization and via load balancing devices. By having access to high-granularity, intra-cabinet data, it is possible to manage server heat loads by first testing new configurations in a virtual setting. Once moved, the effect of load changes can be tested by observing actual energy changes in CRAC units and via total cooling energy consumption.

We propose that the following steps can be taken to achieve substantial energy savings in most data centers:

- Implement meter-grade wattage monitoring for all circuits containing computing and cooling loads
- Establish a baseline data center efficiency using PUE as a measure as recommended by the Green Grid
- Establish heat load profiles for each cabinet by circuit
- Determine which cabinets show the most out-of-balance conditions between circuits
- Provide intra-cabinet balancing for each cabinet with 2 or more circuit-related zones .
- Match cooling resources with actual cabinet heat loads
- Measure the energy effects of moves on nearby CRAC units and in chiller units to see total cooling energy usage
- Re-balance and manage according to cooling responsiveness
- Track PUE on an ongoing basis and respond as necessary

### **About the Authors:**

Robert Hunter is the founder and CEO of TrendPoint Systems. He has over 18 years experience in building, managing and maintaining data centers. Robert was previously the founder and CEO of NetBrowser Communications, the pioneering leader in web-based data center management.

Chet Sandberg is the Technical Strategy Manager for TrendPoint. Chet is the former Chief Scientist of Raychem Corporation, from which he retired in 2002. He received his BS degree from MIT and his Masters from Stanford University. He is a world-recognized leader in the field of heat related sciences. He is an IEEE Fellow, a Senior Member of ISA, a member of ASME, NFPA, and SPE.

## Footnotes

1. Managing the Impact of Increasing Microprocessor Power Consumption, Intel Technology Journal Q1, 2001. Stephen H. Gunther, Desktop Platforms Group, Intel Corp., Frank Binns, Desktop Platforms Group, Intel Corp., Douglas M. Carmean, Desktop Platforms Group, Intel Corp., Jonathan C. Hall, Desktop Platforms Group, Intel Corp.
2. Impact of Process Variations on Multi core Performance Symmetry, Eric Humenay, David Tarjan, Kevin Skadron, Dept. of Computer Science, University of Virginia, February 2007.
3. The Implementation of a 2-Core, Multi-Threaded Itanium Family Processor, IEEE JOURNAL OF SOLID-STATE CIRCUITS, VOL. 41, NO. 1, JANUARY 2006 197, Samuel Naffziger, *Member, IEEE*, Blaine Stackhouse, *Member, IEEE*, Tom Grutkowski, Doug Josephson,, Jayen Desai, Elad Alon, *Senior Member, IEEE*, and Mark Horowitz, *Fellow, IEEE*
4. Full-Chip Analysis of Leakage Power Under Process Variations, Including Spatial Correlations, Hongliang Chang, Dept. of Computer Science and Engineering, University of Minnesota, Sachin S. Sapatnekar, Dept. of Electrical and Computer Engineering, University of Minnesota, *DAC 2005*, June 13–17, 2005, Anaheim, California.
5. Accurate and Efficient Gate-Level Parametric Yield Estimation Considering Power/Performance Correlation, Ashish Srivastava, Saumil Shah, Kanak Agarwal, Dennis Sylvester, David Blaauw, and Steve Director, University of Michigan
6. Green Grid Metrics: Describing Data Center Power Efficiency:, Technical committee White Paper. The Green Grid. February 20, 2007
7. Garday, Doug. May 2007. "Reducing Data Center Energy Consumption with Wet Side Economizers," an Intel white paper
8. Exergy Analysis of Data Center Thermal Management Systems, Amip J. Shah, Van P. Carey, Department of Mechanical Engineering, University of California, Berkeley, CA 94720-1740; Cullen E. Bash, Chandrakant D. Patel, Hewlett Packard Laboratories; Journal of Heat Transfer Copyright © 2008 by ASME FEBRUARY 2008, Vol. 130 / 021401-1
9. M. Norota, H. Hayama, M. Enai, and M. Kishita, "Research on Efficiency of Air Conditioning System for Data Center," *Proceedings of the IEICE/IEEE International Telecommunications Energy Conference (INTELEC)*, October 19–23, 2003, pp. 147–151
10. Thermal Considerations in Cooling Large Scale High Compute Density Data Centers, Chandrakant D. Patel, Ratnesh Sharma, Cullen E. Bash and AbdImonem Beitelmal, Hewlett-Packard Laboratories
11. Thermal Performance Metrics for Arranging Forced Air Cooled Servers in a Data Processing Cabinet, Jeff Rambo, Yogendra Yoshi, Transactions of the ACME, Volume 127, December 2007, pages 452- 459
12. An Approach for the Robust Design of Data Center Server Cabinets, Nathan Rolander, Georgia Institute of Technology, 2005
13. "Temperature's Rising...or Not", Processor Magazine, January 8, 2008, Volume 30, Issue 2, page 19
14. Energy Efficiency in Data Centers: A New Policy Frontier, By Joe Loper and Sara Parr, Alliance to Save Energy with support from Advanced Micro Devices, January 2007
15. Hot Air Isolation Enables Air Cooling in High Density Data Centers, Ian Seaton, Chatsworth Products White Paper, October 15, 2007.
16. High-density heat containment, ASHRAE JOURNAL, December 2007, by Mitch Martin, Mukesh Khattar and Mark Mermagian