

WHITE PAPER

Four Metrics Define Data Center “Greenness”

Enabling users to quantify energy consumption initiatives for environmental sustainability and “bottom line” profitability

By John R. Stanley with Kenneth G. Brill, and Jonathan Koomey, PhD.

Executive Summary

Uptime Institute (*Institute*) research supporting this paper strongly suggests that it is possible for a large-scale enterprise data center to meet business requirements for Information Technology (IT) availability and performance while *profitably* achieving or exceeding many if not all corporate environmental sustainability objectives.

A true *green* data center first meets business requirements while being environmentally, operationally, and financially sustainable. In this paper, the *Institute* develops concepts and specific metrics enabling data center end users to identify profitable efficiency opportunities. Acting on these opportunities can potentially save large amounts of money across a data center portfolio by deferring or possibly eliminating major new capital investment.

In this paper, the term “profitable” means contributing to the bottom line of a business unit or an organization via substantial OpEx/CapEx savings, while still maintaining equal or greater IT performance and capability. The *Institute* firmly believes reliable computing availability and energy/power efficiency are mutually compatible business objectives. Any drive for efficiency cannot lose sight of the business purpose for data centers—information availability for effective decisions and delivering those services effectively to customers with the reliability they require.

Harvesting of these substantial savings initially requires active participation by senior executives (CEO, CFO, CIO, CTO, and EVP Corporate Real Estate), who must break organizational barriers and reset existing perverse incentives that currently inhibit energy and power efficiency within their data centers. Prior *Institute* research suggests that harvesting “five gold nuggets” of efficiency in data centers can in many cases reduce facility energy and power consumption by 30-50 percent with little investment.

This paper defines four factors for determining relative data center greenness. These metrics enable energy and power consumption benchmarking and continuous improvement processes. Three of the four green metrics—*IT systems design and architecture*, *IT hardware asset utilization*, and *IT hardware efficiency*—are under the direct jurisdiction of enterprise IT. The fourth—*site physical infrastructure overhead*—is a shared responsibility of the IT and corporate real estate (CRE)/facilities organizations. These metrics are intended to be intuitively understandable, scientifically accurate, flexible (for inclusion of as yet unknown new technologies), vendor neutral, and contain minimal hidden assumptions and simplifications. Most importantly, the four metrics allow a very complex whole-system problem to be broken down into manageable pieces which can be independently optimized by different parts of an organization.

Breaking down a green data center into IT strategy, asset utilization, hardware efficiency, and site infrastructure overhead metrics enables quantification of current consumption problems

and rapid identification of the highest priority opportunities for improvement. Creating a baseline set of consistent metrics and then improving them will lead to significant savings of both OpEx costs and precious data center UPS and cooling capacity. Senior executives will need to be involved in this process, because current organizational barriers to information flow and perverse incentives must be removed in order to harvest the cost savings. Enterprises failing to take these actions will be at increasing competitive disadvantage relative to more nimble early adopter organizations

This Paper:

- Creates a holistic framework for understanding data center energy efficiency and power efficiency, to help executives understand how to get the most from data center investments
- Proposes what it means to be a *green* data center, enabling IT organizations to tie their information availability and operational/financial sustainability objectives into their firms’ broader environmental sustainability goals (particularly carbon footprint reductions)
- Advances the current conversation for how standard data center energy efficiency metrics should be defined, quantified, and measured, so that industry stakeholders and policy makers can align around broadly acceptable metrics
- Uses terms that are scientifically accurate, reasonably intuitive, and motivate *whole-systems* optimization
- Breaks overall data center energy efficiency into four manageable chunks, so that data center decision-makers can understand their options, and so that the providers of IT hardware and facilities infrastructure can see the many different ways of meeting their customers’ needs
- Delegates responsibility for each metric to specific parties within the data center, so that profit-enhancing energy efficiency improvements can move forward with clear measurement and accountability by organizational function

Table of Contents:

Background	3
Four Green Metrics, Four Accountable Functions	4
Scope of This Paper	6
Defining the Green Data Center	7
Criteria for Good Metrics	8
<i>Power and Energy are Different</i>	8
Where to Measure Data Center Power and Energy Consumption	10
<i>Whole-Systems Thinking: Performance vs. Productivity</i>	10
Additional Metrics Details	11
<i>Institute Definitions</i>	12
Data Center Power Consumption “at the Meter”	12
Hardware Load at the Plug	12
Hardware Compute Load	13
Site Infrastructure Power Overhead Multiplier (SI-POM)	13
IT Hardware Power Overhead Multiplier (H-POM)	14
DC Hardware Compute Load per Unit of Computing Work Done	14
Deployed Hardware Utilization Ratio (DH-UR)	15
Deployed Hardware Utilization Efficiency (DH-UE)	15
Free Cooling	16
Enable Energy Saving Features	18
Putting it All Together	19
Conclusion and Next Steps	19
References	20
White Paper Reviewers	21
About the Authors	21
About the Uptime Institute, Inc.	21

Background

Data center electricity consumption has been getting increased attention and concern from both industry leaders and public policy makers¹. Business leaders are now becoming aware of the increasing costs of electricity and of the site infrastructure necessary to support each new piece of IT equipment they deploy². Both corporate leaders and policy makers are increasingly aware of the implications of the capacity limitations of national and regional power grids and of carbon emissions attributable to rapidly growing data center electricity use.

Many agree that substantial benefits to both business and environmental interests accrue from making data centers significantly more energy efficient. Most would also agree that poorly defined or poorly implemented efficiency metrics could impede business innovation while still failing to achieve environmental sustainability goals.

What does *efficiency* really mean in the context of a data center? How is it to be measured, quantified, and broken into chunks small enough for different organizations within the corporate enterprise to take directed action? In this paper, the *Institute* intends to answer these questions in a way that advances both the intellectual and the practical discussion for organizational change among users as well as all other industry stakeholders.

Much work has already been done to define and quantify data center efficiency. As early as 2003, the *Institute* promulgated a research-based measurement methodology. Introduced in client work as the Triton Coefficient of Effectiveness, this was defined as the total utility energy required to operate a data center divided by the critical electrical load in the computer room.

At the same time, the *Institute* began discussion with its 100-corporate member Site Uptime Network[®] on the need to measure the ratio of data center utility meter load to UPS output. First introduced at its Symposium 2006 on high-density computing, the *Institute* began calling this metric the Site Infrastructure Energy Efficiency Ratio (SI-EER), which measures the number of kW needed at a data center's outside utility meter to deliver one kW of reliable power to the IT equipment inside. The *Institute* also developed the Information Technology Energy Efficiency Ratio (IT-EER), which measures the compute performance of IT equipment per embedded Watt of power consumption. Combining these two terms together yields the overall Data Center Energy Efficiency Ratio (DC-EER)³.

The Power Usage Effectiveness (PUE) metric was introduced by The Green Grid in 2007. PUE is defined as the total data center electrical load divided by the IT electrical load. Authors Malone and Belady also defined the reciprocal of PUE, called Data Center Efficiency (DCE), as well as an integrated metric called Data Center Performance Efficiency (DCPE), which is a measure of the ratio of useful work done in a data center to total facility power⁴.

Metrics and benchmarking methods are also being discussed by public policy makers, in particular within the EPA's Energy Star program. The recently published ENERGY STAR *Report to Congress on Server and Data Center Energy Efficiency*⁵ includes a discussion of current government and industry work on metrics. The ENERGY STAR program is currently exploring measurement protocols and voluntary efficiency specifications for servers and possibly other data center IT and infrastructure hardware⁶.

The ENERGY STAR report also references other current work on efficiency metrics, including a technical paper on server energy measurement protocols by Dr. Jonathan Koomey (2006), work done by the SPEC Power and Performance Committee (2007), and a white paper authored by Neil Rasmussen of APC (2006). (See the *References* section at the end of this paper for more information on these three efforts and other reference materials.)

Four Green Metrics, Four Accountable Functions

The *Institute* groups energy and power efficiency metrics for a green data center into four broad categories, each of which is addressable by a different function within a user organization.

Too often, organizations say that “efficiency is everyone’s job.” What statements like that really mean is that “it’s no one’s job.” Delegating the responsibility for each category of efficiency potential to a specific group within a data center organization creates accountability for ensuring that gap analysis between current practices and best practices actually gets done and that beneficial efficiency improvements actually are made.

In addition, categorizing the metrics in this way breaks the formidable task of data center efficiency into manageable chunks, giving accountability leaders and teams a place to begin. The four metric categories are:

- **IT strategy (business requirements, systems architecture and platform selection, data topology, and network design)**—CIOs/CTOs should ask themselves, *How can we achieve our business objectives with less energy by considering different IT design, architecture, and hardware options to achieve computing and network availability, reliability, and performance?*

This factor is the responsibility of the CTO/CIO, who occupies Quadrant 1 (Q1), in the *Institute's* Integrated Critical Environment™ (ICE) Team organization. (See below for further ICE Team definition in text and in **Figure 1**). Examples of these options could include alternatives to existing high-availability strategies like active-active data centers, use of mainframes instead of servers, utilization of different data center site Tier functionality levels for different application availability requirements, centralized versus decentralized processing, number of redundant global data set copies, disaster recovery, etc.

1 See *Opinion of the Institute: EPA Report Should Spur Industry-Wide Green Data Center Movement*, <http://www.uptimeinstitute.org/resources>

2 For a more thorough discussion of these costs and what they can mean for IT-intensive businesses, see the *Institute's* recent white paper entitled *The Invisible Crisis in the Data Center: The Economic Meltdown of Moore's Law*, <http://www.uptimeinstitute.org/resources>

3 For the early use of these terms, see the *Institute's* March 2006 paper *High Density Computing: The Path Forward 2006*, <http://www.uptimeinstitute.org/resources>. In January 2007, the *Institute* referred to similar concepts in *Data Center Energy Efficiency and Productivity*, <http://www.uptimeinstitute.org/resources>, although the names of the terms had changed slightly.

4 For the complete description of PUE and related metrics, see The Green Grid's February 2007 publication, “Green Grid Metrics: Describing Datacenter Power Efficiency.” Note, however, that the concept of PUE was introduced much earlier, by C. Malone and Christian Belady in September 2006 at the Digital Power Forum.

5 The report can be found on ENERGY STAR's website at: <http://www.energystar.gov/datacenters>

6 See Section 7.2.2 of the EPA report.

- **IT hardware⁷ asset utilization**—CIOs, CTOs, and senior data center executives should ask themselves, *How can I maximize the fraction of my IT hardware assets which are deployed productively and fully utilized?*

This factor is the responsibility of the data center management function (a combination of Quadrants 1 and 2 in the *Institute’s* ICE Team organization). Examples of efficiency options here include turning off comatose servers and storage, appropriately virtualizing both applications and storage, and enabling server power-save features⁸.

- **IT energy and power efficient hardware deployment**—IT architecture and capacity planning executives and purchasing managers should ask themselves, *How can I select and justify buying IT hardware that delivers the most effective computing performance per Watt of power consumption at the plug?*

This factor is a joint responsibility of the IT architecture and capacity planning function along with the specifying and procurement function (Q1 in the *Institute’s* ICE Team organization). Hardware can be made more efficient by utilizing better power supplies which more efficiently convert AC power at the plug to DC power used by the internal components, or by selecting internal components that do more useful work with less DC power (more efficient chips, drives, etc.)

- **Site physical infrastructure overhead**—Facilities or corporate real estate (CRE) decision makers should ask themselves, *How can I maximize the amount of useful power/energy delivered to the power plugs of IT hardware for each unit of power/energy consumed at the data center utility meter and thereby reduce site infrastructure “overhead?”*

This factor is primarily the responsibility of the facility management function with integral support from data center operations (Q3 and Q2 in the *Institute’s* ICE Team organization) to design and operate the site physical infrastructure efficiently while still achieving business appropriate levels of uptime, concurrent maintainability, and fault tolerance.

The first three metrics combine or nest into a composite IT productivity per Watt of power consumed by the IT hardware. The word “productivity” has been intentionally chosen to convey that what counts in a green data center is the useful computational work performed⁹. The final category of site physical infrastructure overhead converts IT power consumption to data center utility power consumption. These ideas are further developed later in this paper.

Despite the fact that these categories are delegated to different parties, it is essential to maintain good, open communications and information sharing, especially across the divide that typically

Figure 1—Integrated Critical Environment Teams

	Site Infrastructure Overhead	IT Productivity Per Watt
Strategic	Quadrant 4 (Q4): EVP Corp Real Estate <ul style="list-style-type: none"> • Benchmark SI-EOM with peers • Manage Corporate Social Responsibility • Manage the organization’s global data center asset portfolio 	Quadrant 1 (Q1): CIO/CTO <ul style="list-style-type: none"> • Benchmark existing hardware asset utilization and harvest lowest hanging fruit • Change IT governance to better allocate resources • Incorporate energy efficiency in hardware selection decisions • Benchmark IT Productivity Per Watt with peer organization • Incorporate energy considerations into IT Strategy
Operations	Quadrant 3 (Q3): Facility Manager <ul style="list-style-type: none"> • Measure and manage SI-EOM • Implement known SI Overhead best practices • Utilize more free-cooling • Install more energy efficient infrastructure components (when possible) 	Quadrant 2 (Q2): Data Center Manager <ul style="list-style-type: none"> • Implement Q1 harvesting decisions • Implement known computer room best practices

⁷ Software utilization is also a factor, but not explicitly considered here. Operating system and application bloatware requires larger and more powerful servers. Consolidation of multiple ERP systems serving different functions to a single system on a single platform will reduce total power consumption in addition to other benefits.

⁸ “Comatose” servers are those which are installed and running, but which are running applications no longer needed or which run no applications at all.

⁹ There are many ways to increase the amount of useful work done per unit of power or energy consumed. For example, turning off comatose servers saves power while sacrificing only compute cycles that were not doing anything productive anyway.

exists in most organizations between the IT and CRE/facilities divisions. The groups must foster close cooperation to evaluate various tradeoffs and determine optimal solutions when viewed at the level of the entire data center.

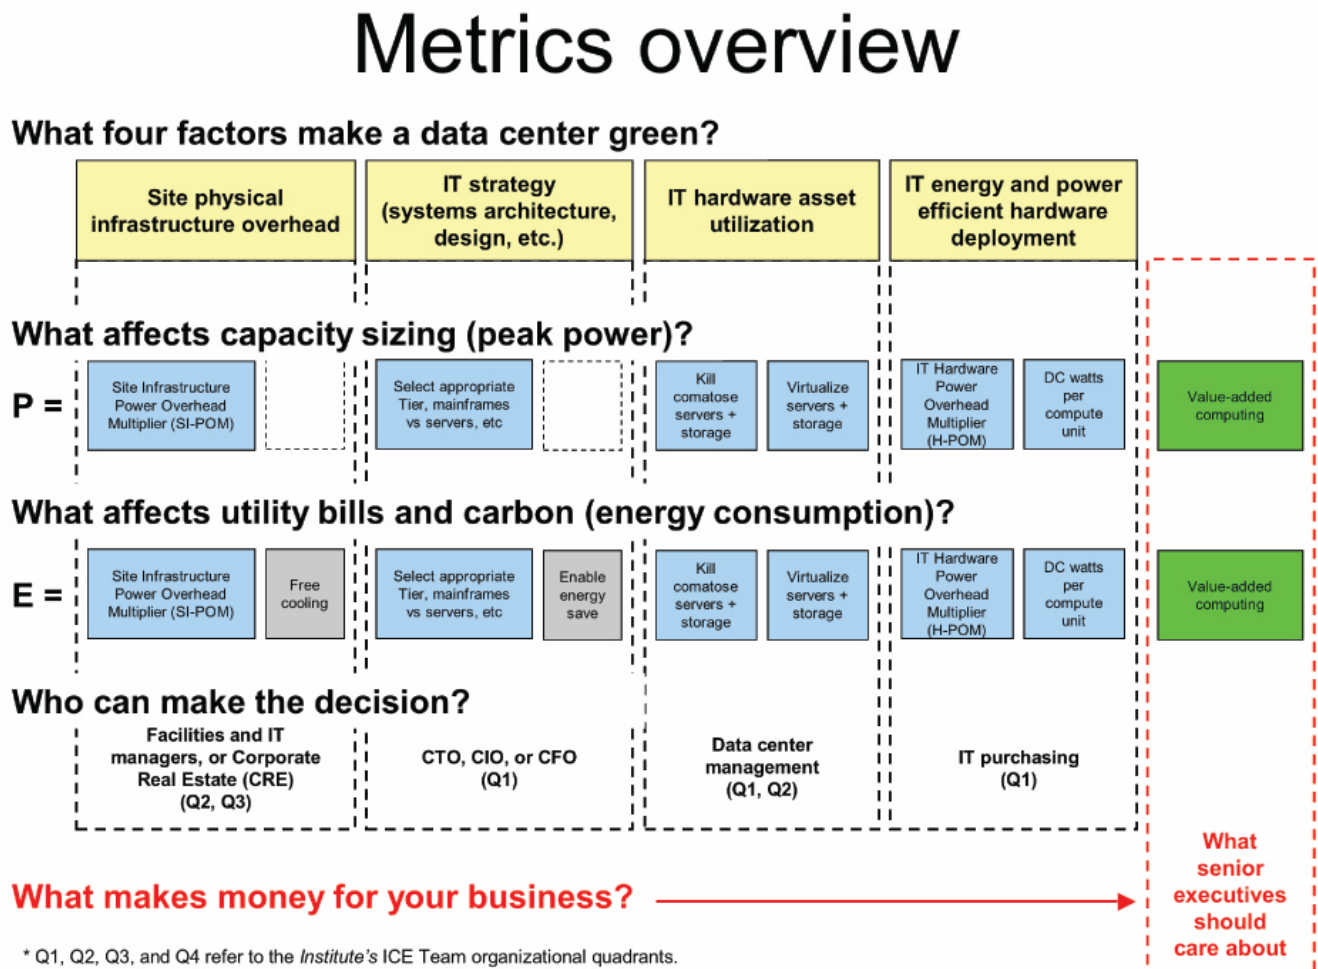
To meet this need, the *Institute* has developed a management protocol and methodology and recommends that organizations create four-quadrant Integrated Critical Environment (ICE) Teams with seasoned, well-respected, decisive leadership having specific executive purview for energy/power reduction goals and timetables¹⁰. To overcome existing organizational barriers, information silos, and inertia, these teams must be chartered by the corporate executive committee and report back their progress at least semi-annually (perhaps even quarterly at the outset). Only then can these teams operate effectively enough to realize the potential cost savings from improved efficiency.

Figure 2 is a diagram of data center energy and power efficiency categories and the four green metrics into which the *Institute* has grouped them. Note that three of these factors are completely within the IT jurisdiction, and that the fourth, while more closely identified with the facilities or CRE departments, still has a strong IT relationship. No longer is energy efficiency something to be addressed only by facilities/CRE personnel. Note also that each of the four categories can contain multiple sub-metrics for different efficiency opportunities which can be separately benchmarked and individually addressed.

Scope of This Paper

In presenting this paper, the *Institute* intends that its primarily large-scale data center user-experience-oriented perspective can further advance and refine the dialogue on energy efficiency within the data center community. This paper intends to complement, not contradict, previous work. Ultimately, the

Figure 2—Overview of Green Data Center Metrics



¹⁰ For more information on ICE Teams, see the *Institute's* 2007 white paper *Data Center Energy Efficiency and Productivity*, <http://www.uptimeinstitute.org/resources>. See also the upcoming *Institute* white paper on chartering, empowering, managing, and measuring ICE Team performance.

goal is for all stakeholders, especially users who have the most to benefit, to unite behind a comprehensive set of data center efficiency metrics that boost user organizational performance, encourage vendor innovation, enhance IT productivity, and advance broader energy and environmental goals.

The *Institute* is committed to developing a comprehensive set of metrics which look at the data center as a *whole system*. By this, the *Institute* means not just the data center building and its mechanical and electrical systems (which are the only factors currently considered in most efficiency metrics), but the energy productivity and utilization of IT assets within the building, which are ultimately far more significant. Users need ways to prioritize data center energy initiatives to focus on where the biggest *enterprise* savings are irrespective of organizational boundaries.

To see why it's important to look at what's inside the data center as well, consider the following example. *Institute* research suggests that one of the biggest and quickest initial data center efficiency improvements can come from IT simply turning off "comatose" servers. As mentioned in the footnote above, comatose servers are those which are installed and running, but which are running applications no longer needed or which run no applications at all. Such servers have been functionally replaced, but they can account for 15 to 30 percent of the total servers installed in some data centers, because various kinds of organizational barriers and perverse incentives¹¹ prevent them from being identified, turned off, and removed from service. Simply turning off comatose servers can save large amounts of OpEx and CapEx at little or no cost¹². Even better, deactivating these servers is something that IT managers can do *immediately*. Contrast this opportunity with improvements to an inefficient data center's building shell or cooling system, which can be virtually impossible to achieve because data centers typically can't be taken off-line for renovations. This example thus highlights the need for metrics which take a holistic look at the actual computational work being performed inside a data center *in addition* to the architectural elements of the data center's building shell and mechanical and electrical (M&E) system components therein. Including the energy efficiency of the computational processes inside the data center is a radical departure from previous metrics work which has primarily been limited to the building shell and M&E systems.

Readers should also note that the scope of this paper does *not* include the detailed technical procedures for actually measuring any of the four proposed metrics. For example, Koomey, *et al.* (2006) describes the many variables that must be standardized to perform accurate hardware efficiency benchmarking of a server. These include the exact input power voltages used, the resolution of the measurement equipment, the specific compute jobs tested, and the ambient temperature and relative humidity at which measurements are performed.

In this paper the *Institute* seeks to foster an industry discussion on energy and power efficiency metrics at a *meta* level. After such meta-layer understanding and agreement is reached, the *Institute* and others can refine the important technical details of the four individual metrics' measurement. Refining these technical details will not be easy by any means¹³. However, at least these subsequent efforts can be put forth within the context of the "big picture" upon which there is broad community consensus.

Defining the Green Data Center

Data centers have a variety of environmental impacts, but those relating to their substantial electricity use are by far the most important. The electricity intensity of typical data centers (in Watts/ft²) is 10 to 20X greater than that of office buildings, and it continues to grow rapidly¹⁴. For this reason, energy use during operation represents a far greater share of a data center's environmental impact than other characteristics such as energy embodied in building materials, impacts on the immediately surrounding environment (particularly during construction), and choices on insulation, roofing, and M&E components. (One exception to electricity consumption being the dominant data center environmental concern is the quantity of make-up water used in cooling towers for evaporative heat, which can be in the tens of thousands of gallons per day.)

Senior business leaders (some at the urging of Boards of Directors) are now beginning to wonder how data centers might fit into the overall environmental sustainability objectives of their firms, particularly those objectives related to climate change. Data centers themselves are not a primary source of greenhouse gas emissions, and much of the computational work done in data centers may reduce emissions in other areas of an organization¹⁵. However, since the electricity which data centers consume typically comes from fossil-fuel-fired power plants, large and rapidly growing levels of data center energy use are a matter of corporate, public policy, and regulatory concern.

For these reasons, the *Institute* is focusing our current and future green data center initiatives on defining, measuring, and managing the four key performance indicators (KPIs) of energy and power efficiency for both IT hardware and site physical infrastructure. In addition to being the most important, these meta efficiency KPIs are also the most technically difficult to define. These meta KPIs create a conceptual umbrella under which additional sub-KPIs will evolve over time.

Other thought leaders in the building shell efficiency space (for example, the US Green Building Council's LEED program¹⁶) may wish to create more specific measures of environmental sustainability for a data center's site and building envelope that the *Institute's* proposed four metrics do not currently address (materials sourcing, use of toxic materials, recycling of retired IT equipment, etc.). The *Institute's* goal in this paper is to stimulate

¹¹ One example of a broken incentive structure is the fact that in many companies, users or IT managers who specify servers typically do not directly pay for the supporting facilities capital (UPS, cooling systems, etc.) or for the electricity to operate the equipment. These costs are the responsibility of a separate facilities organization, so IT managers often have little or no incentive to specify more efficient servers. Worse, since it's the IT managers who decide what to order from hardware vendors, the vendors have little incentive to consider efficiency in their equipment designs.

¹² CapEx for a future data center capacity addition is saved or deferred because each comatose server removed frees up precious UPS capacity which can be used later when additional productive servers need to be installed.

¹³ For example, since different types of server configurations excel at different types of compute jobs, selection of particular compute routines for benchmarking is likely to be quite contentious. SPEC has made much progress in this particular area, but similar technical disagreements will no doubt arise around many of the metrics.

¹⁴ Data centers have such a high power density because there are lots of power-consuming servers, storage devices, and networking devices stacked into a relatively small floor area.

¹⁵ For example, logistics calculations can reduce necessary shipping miles. IT also enables telecommuting, video conferencing, and other travel-reducing work options. Atkinson and McKay discuss some of the economic and environmental benefits of IT in their publication *Digital Prosperity: Understanding the Economic Benefits of the Information Technology Revolution*

¹⁶ Leadership in Energy and Environmental Design (LEED) is a nationally recognized green building rating system developed by the US Green Building Council, <http://www.usgbc.org/LEED/>

discussion within and among IT users, all industry stakeholders that provide products and services to users, and public policy makers. In the past, such discussions failed to include the entire community to ensure that everyone would understand the technical and economic consequences of decisions and that all dissenting voices were heard. This conversation needs to be high-level and sufficiently comprehensive to provide an actionable list of the important components of overall data center power and energy consumption. When that goal is reached, data center operators who successfully address each of the four efficiency improvement areas and who have instituted a comprehensive and continuous improvement program can be said to be green in the sense that they have taken advantage of most or all major opportunities to reduce their environmental impact.

Furthermore, the *Institute* believes that the greening of data centers via improved energy and power efficiency will prove to be sufficiently cost-saving (i.e., directly contributing to bottom line profits) that most organizations will aggressively pursue *efficiency-for-profit* on their own, in the absence of public policy mandates or regulatory threat.

However, the authors still see a substantial *public education and opinion-influencing role* for policy makers in the data center efficiency discussion, given their considerable knowledge of both efficiency opportunities and environmental goals. Lack of “C-suite” awareness about the need to reset organizational barriers and overcome perverse incentives continues to be a major inhibitor to harvesting promising efficiency opportunities.

Criteria for Good Metrics

Given the considerable contributions that improved data center efficiency can make to both organizational profits and environmental goals, good metrics for evaluating energy and power efficiency are essential. To make future stakeholder discussions about metrics as productive as possible, the *Institute* proposes these criteria below for judging the quality of metrics as they evolve.

Effective metrics should:

- Be intuitive:
 - » The meaning should be as clear as possible, just from the name.
 - » It should be obvious whether it is good for a metric to go up or down¹⁷.
 - » It should be simple as possible, but no simpler. (*i.e., metrics should not gloss over important details.*)
- Be scientifically accurate and used precisely.
- Be granular enough that individual aspects can be analyzed in manageable chunks and assigned to specific parties for improvement.
- Nest together elegantly so that the overall equation for data center power or energy use is simple.
- Make it easy to calculate how much of a data center’s power or energy is wasted by a particular practice, process, or configuration.

- Work for any type of IT hardware—servers, storage, networking, and future products not yet invented.
- Be sufficiently flexible to respond to new technology developments. (*Metrics should not need to be scrapped and rewritten each time the industry develops a new, innovative way to increase efficiency.*)
- Clearly distinguish between power and energy. (*These are technically and semantically distinct and affect different aspects of data center operation.*)
- Be precise to avoid ambiguous terms. (*i.e., does “full load” mean “full compute load” or “full power load?”*)
- Be vendor-neutral¹⁸ without hindering vendor-dependent R&D innovation.

The *Institute* invites further discussion of these criteria, in addition to the inquiry into the specific green metrics themselves.

Power and Energy are Different

Distinguishing between the terms *power* and *energy* is extremely helpful to the metrics discussion because it forces efficiency advocates to be clear about what kind of benefits a particular type of efficiency improvement actually offers. Some improvements save only power, some save only energy, and most save some of both.

In a data center, peak *power* demand determines how big the UPS, cooling system, and utility feeders must be. Saving power tends to reduce the CapEx investment for new data centers and allows deferral of expensive capacity expansions to existing data centers. These capital savings have much greater economic and profitability impacts on data center operators than does the utility bill, an OpEx item. Thus, the majority of data center operators will be most interested in saving power, which reduces or defers CapEx.

On the other hand, policy makers and corporate level enterprise sustainability initiatives are typically more interested in *energy*, because energy consumption determines the quantity of fuel which must be burned at a power plant somewhere, with its requisite emissions of carbon dioxide and other pollutants¹⁹.

Here are some useful notes regarding the terms:

- *Power* is a spot measurement at a particular point in time.
- *Energy* is consumption over a *period* of time.
- Power is measured in kilowatts (kW), while energy is measured in kilowatt hours (kWh)²⁰.
- Although there is a power measurement for every instant, the measurement that users typically care about is the “peak” or highest power draw at any point in time.

These distinctions are important because it’s possible for systems to draw a lot of power without using much energy, and vice versa. For example, a 1,500 watt (W) hair dryer that is used five minutes each day uses 15X more instantaneous power when running than

¹⁷ For example, the *Institute’s* original Site Infrastructure Energy Efficiency Ratio (SI-EER) fails this test. Recall that SI-EER is the quotient (kW of power delivered at the utility meter) / (kW of power that makes it to the IT equipment). It sounds good when the “Efficiency Ratio” is going up, but that means *less* of the site’s power is making it to the IT equipment, which is bad!

¹⁸ Vendor neutrality should be an obvious necessity, but the authors specifically include it here for completeness.

¹⁹ The function within facilities that pays the utility bills will also most likely be interested in energy consumption, but often this group is totally isolated from IT and data center operations.

²⁰ Some data centers use steam or natural gas for energy. The reporting method for these sites will need to be adjusted to account for their alternative energy source(s).

a 100W light bulb left on all year, but the hair dryer uses only about 1/20 as much energy²¹. An IT manager who hibernates idle servers at night will save *energy*, but not reduce peak *power* demand. The site physical infrastructure system’s capacity must still be sized to accommodate times of peak compute demand when all servers are running their hardest.

Unfortunately, many analysts and even some energy experts inaccurately use the terms power and energy interchangeably. The *Institute* recommends that future discussions of data center energy/power efficiency be technically precise about which term is meant.

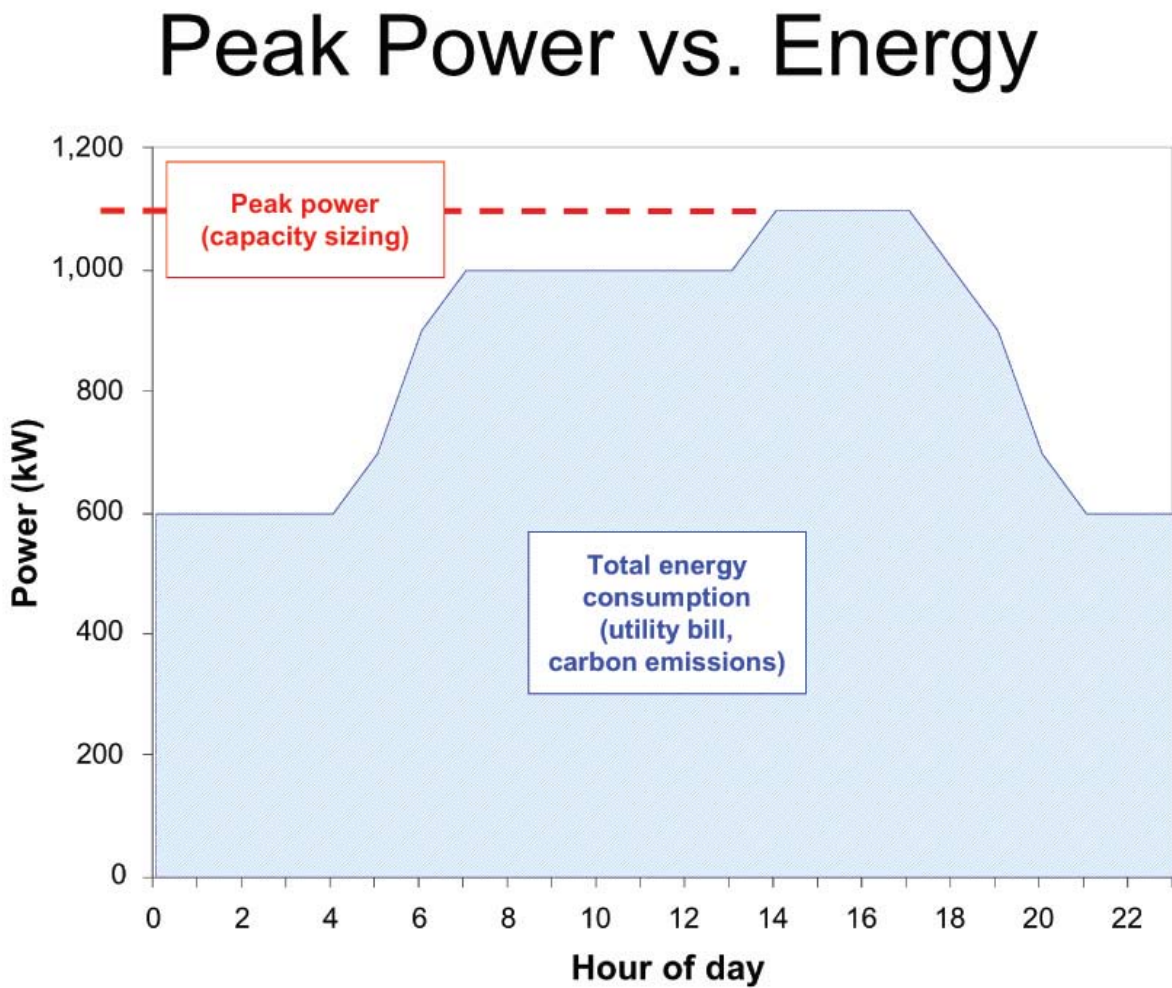
If a data center really draws the same amount of power every hour of the year, then the energy use in kWh is just the power in kW

times 8,760 hours. However, in a green data center, annual energy use will not be the peak power times 8,760 hours, because:

- Idle servers and storage can be set to hibernate at off-peak times, at night, or on weekends,
- Free cooling can be used during some hours of the year, and
- Future generations of IT equipment might be designed to draw less than their peak power when they’re only running at partial compute load.

Figure 3 provides a graphical representation of the relationship between power and energy. **Table 1** compares power and energy in terms relevant to data center operators.

Figure 3—The Relationship Between Power and Energy



²¹ Hair dryer = (1,500 watts) x (5 minutes) / (60 minutes per hour) x (365 days) / (1,000 watts per kW) = 46 kWh. Light bulb = (100 watts) x (8,760 hours per year) / (1,000 watts per kW) = 876 kWh.

Table 1—Power and Energy Comparison

Power	Energy
kW	kWh
CapEx	OpEx
Data center system capacity constraints	The electricity bill
Demand charge / utility feeders	kWh charge
--	CO ₂ emitted
Spot measurement, typically at peak time	Sum, typically over the entire year
Full load efficiency	Part load efficiencies
Site Infrastructure Power Overhead Multiplier (SI-POM) Hardware Power Overhead Multiplier (H-POM) Compute productivity per Watt of internal power Virtualization Kill comatose servers and storage	
--	Free cooling Enable energy save feature

Where to Measure Data Center Power and Energy Consumption

Measuring power and energy consumption in a data center occurs in three important places. These will be familiar to those who have examined previous literature on metrics.

The first is total data center consumption *at the utility meter*, which determines the utility bill (both the energy portion and the power-based peak demand charge). Whether a utility can physically meet the demands of the site without having to increase the capacity of the power feed is also determined by peak *power* consumption at the utility meter.

The second place to measure power and energy consumption is *at the plug* bringing conditioned AC input power to the input of the IT equipment²².

In most data centers, about half the power and energy consumed at the utility meter never makes it to the IT equipment; instead it's consumed to run cooling equipment or dissipated in conversion within the UPS and PDU systems. The *Institute* calls this the site infrastructure *overhead*. The objective over time for all stakeholders should be to drive such losses down.

The third place to measure is the *hardware compute load* inside the box of the IT equipment itself, to see how much of the power and energy at the plug actually makes it to the internal components which do the useful computational work. A significant portion of the power and energy that go into the box are either wasted in power supply conversion losses or diverted to non-computing loads such as internal fans for rejecting heat.

Figure 4 shows a schematic of data center electricity flow and the relevant measurement locations.

Whole-Systems Thinking: Performance vs. Productivity

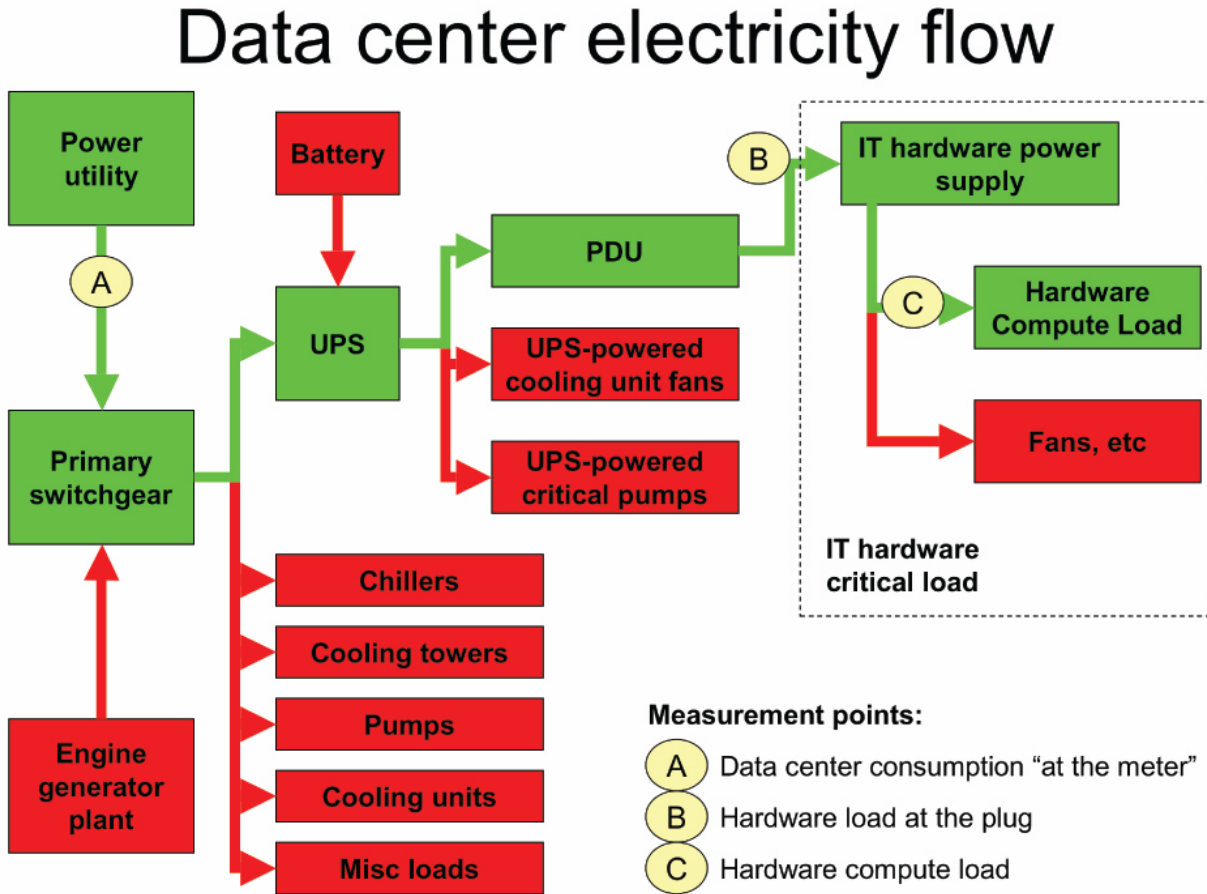
Data centers must consider both computational and network availability needs and environmental goals. As these apparently conflicting goals are considered, decision makers will need to think critically about *what* strategies to implement in addition to their regular thinking about *how* best to implement them.

For example, data center cooling system design engineers and facilities managers may already be accustomed to searching for chiller equipment that uses the fewest possible kW per ton to produce chilled water. Clearly, doing this will improve cooling efficiency. However, these facilities experts may not think to ask *why* the chilled water needs to be at the 42°F (6°C) commonly specified for an office building instead of, say, 49°F (9°C). From a whole-systems perspective, standard efficiency chillers at 49°F (9°C) would deliver cooling more efficiently than premium efficiency chillers at a lower set point, because latent dehumidification in the data center is eliminated. Similar initiatives can be employed to reduce the volume of outside air introduced into computer rooms using CO₂ sensors to satisfy code requirements while also being energy efficient.

Analogous examples exist on the IT hardware side for servers and storage. Rather than simply find the most efficient 7,200 rpm storage drive, astute storage managers could well ask why infrequently-used data needs to reside on such a fast device. Power consumption goes up as the cube of spindle speed, so 7,200 rpm drives consume 8X more power than 3,600 rpm drives. Perhaps rarely-used data could be stored on slower drives or even on tape, both of which would consume less power and energy. This change may be economically implemented without necessarily buying new or more hardware. Similarly, business-line application managers might choose to consolidate 10 or more servers via

²² For data centers using direct current power distribution, users would be measuring conditioned DC input power.

Figure 4—Three Places to Measure Power and Energy Consumption



virtualization during the next IT hardware refresh cycle, rather than simply swapping each box with a new one containing a more efficient power supply. In this example, virtualization immediately cuts power and energy by 90 percent. In sharp contrast, implementing higher efficiency power supplies for all hardware probably costs more, must wait to coincide with a technology refresh, and then may only save five to ten percent of total power and energy.

Clearly, any energy efficiency initiative should not be allowed to reduce IT availability or the quality of service that the data center provides. However, data center decision makers will likely find often-overlooked efficiency opportunities if they remember to consider what new configurations will better serve their firms’ business needs, rather than just repeating previous practice in a slightly better way.

Additional Metrics Details

The remainder of this paper goes into more precise detail about how each of the data center efficiency metrics and sub-metrics is defined. As previously noted, these metrics are intended to help decision makers divide overall data center efficiency into technically accurate and manageable chunks that can be understood, measured, and clearly delegated to appropriate accountable groups for improvement. These metrics are also intended to advance user and industry stakeholder discussions one step closer to general agreement on the meta-level definitions of *efficiency* and *green* as related to data centers. The *Institute* looks forward to receiving feedback and to the ensuing all-stakeholder dialogue.

After clarifying definitions and several measurement points for power/energy, the *Institute* describes the efficiency metrics themselves. Each metrics section includes a basic explanation of whether the metric applies to power, energy, or both; a brief numerical example; and a more detailed technical note. The

technical notes intend to qualify the metrics by discussing their complexities, to explicitly state the simplifying assumptions that have been made in the metric definitions, and to provide the rationale for those simplifications.

Institute Definitions

- **IT hardware (also equipment)**—In this paper, the terms refer to servers, storage devices, and networking devices. Storage devices include both disk and tape, while networking devices include switches, routers, modems, and point-of-presence (POP) equipment.
- **Site physical infrastructure**—This refers to the data center’s cooling equipment, power conditioning equipment, and the building shell necessary to support the IT hardware inside. Readers are cautioned not to confuse site physical *infrastructure* with the common IT definition of infrastructure, which typically only includes IT requirements like racks, power cables, data cabling, POP, etc.
- **Compute**—The *Institute* uses this term to generically refer to the useful work done by the IT hardware. For servers, compute might mean calculating results or processing transactions. For storage devices, computing refers to storing and retrieving information. For networking devices, computing might be routing data packets, analyzing traffic to detect network attacks, etc.

Data Center Power Consumption “at the Meter” (kW_{AC})

This is the total power consumption for the data center as measured at the utility meter or building sub-meter²³, which tells facilities/CRE personnel and utility-demand planners how much power needs to be brought onto the site. This includes *everything* the data center uses, even non-critical loads such as lights. (See **Figure 4, point A**)

Peak at-the-meter power consumption is what determines the demand charge the data center must pay to the power utility provider. Utility companies typically levy demand charges based on the highest instantaneous power draw experienced by a facility in a single month because that, in aggregate for the grid, determines how much total generating capacity will be required. In the past, peak and average demand tended to be very similar so few data centers incurred demand charges. As hardware power-save features are enabled, green data center power consumption will vary with compute load during the day and especially on weekends. As data center loads begin to vary, some facilities may have trouble maintaining the stability of their cooling systems, if that load variance is too wide. Demand charges could also become much more financially significant relative to the energy-based (kWh-based) part of the utility bill.

Some utilities are beginning to offer time-of-use pricing to provide an incentive for users to reduce power consumption during peak-demand periods or during a utility crisis. Some data centers might benefit from such incentives by leaving the grid and running on their own internal power-generating plant (although the *Institute* does not recommend this as a standard practice because long term use has the potential for degrading electrical reliability of the engine generating plant), or by ramping down chillers and using stored chilled water produced during non-peak rate periods as a thermal battery. (The *Institute* whole heartedly recommends this practice because it enhances cooling reliability at all times while saving energy during some periods.)

Energy (kWh):

It’s also possible to measure or compute energy use at the utility meter

Technical note: Strictly speaking, annual energy consumption at the meter will not simply be the peak power consumption at the meter times 8,760 hours, because it’s possible for power load to vary throughout the day or week as IT hardware adjusts to different compute loads and the air conditioning system operates more or less efficiently in response to ambient conditions. Technically, annual energy use in the data center can be defined as:

$$\text{Energy} = \int_{8760 \text{ hours}} P(t) dt$$

where $P(t)$ is the power drawn at a specific instant.

Until large scale implementation of power save features occurs, multiplying peak power consumption at the meter times 8,760 hours is not a bad approximation of annual data center energy use. Although compute load varies throughout the day, many pieces of IT equipment can use as much as 90 percent of their peak power even when idle, so the IT hardware power draw in many cases is relatively flat. Often times, the change in air conditioning efficiency over the range of outdoor weather conditions can be safely ignored, too. Note that if free cooling is being used, the saved energy will need to be subtracted from the kWh total.

In the future, broad scale implementation of existing power save features (similar to those found in laptops) plus improved technology will allow IT equipment to more effectively reduce its power consumption when full computing capability is not needed, so that power demand can be significantly lower during some periods of the day than it is during peak compute periods

Hardware Load at the Plug (kW_{AC})

For a single piece of IT equipment, this is the alternating current (AC) power consumption measured at the hardware power plug²⁴. Summed over all IT equipment in the data center’s raised floor area, this metric tells facility managers how much power must be delivered by the UPS and PDU systems. (See **Figure 4, point B**)

²³ This measurement is easiest to make at stand-alone facilities. At sites where the data center space is part of a larger building, the data center load may not be metered separately. If it is not separately metered, it will be difficult to determine what portion of the consumed electricity is attributable to the data center.

²⁴ The authors acknowledge that in Telco central offices and data centers with direct current power distribution, this would be DC power consumption at the plug.

Energy (kWh):

It’s also possible to measure or compute energy use at the hardware plug, either for a single piece of IT equipment or for an entire data center.

Technical note: As with energy consumption at the meter, energy consumption at the plug is not simply peak power times 8,760 hours, although this can often be a reasonable approximation.

Hardware Compute Load (kW_{DC})

This quantity describes the number of Watts of direct current (DC²⁵) power that are consumed by the computing components within the IT equipment. This power comes from the AC hardware load at the plug and is the net amount left after subtracting AC to bulk DC conversion losses in the hardware’s power supply(s) and power diverted to run the device’s internal cooling fans. (See **Figure 4, point C**).

This quantity is rarely directly observed by users. However, the ratio of AC Hardware Load at the plug divided by DC Hardware Compute Load is a useful metric for manufacturer hardware design teams to consider as they seek ways to provide their customers with more compute performance for less power. The lower this overhead number is, the more efficient the resulting system will be.

For an entire data center, AC Hardware Compute Load is the sum of the individual AC Hardware Compute Loads of all the pieces of IT equipment within it.

Energy (kWh):

It is also possible to compute energy use at the level of AC Hardware Compute Load.

Technical note: Although DC hardware compute load Watts are defined specifically to be “downstream” of losses in the main device power supply, they are still “upstream” of losses in other power conversions such as voltage regulator modules (VRMs) and power lost due to resistance in the circuit boards.

Although these losses can be on the same order of magnitude as the losses in the main device power supply, measuring them is very difficult in both theory and practice (Calwell 2007). For simplicity, the Institute defines Hardware Compute Load Watts to be measured just downstream of the main power supply and fans. DC to DC voltage conversion losses further downstream can be addressed by hardware vendors as they design their products to provide maximum compute performance for minimum power.

As with energy consumption at the meter, energy consumption at the level of Hardware Compute Load is not simply peak power times 8,760 hours, although this can often be a reasonable approximation.

Site Infrastructure Power Overhead Multiplier (SI-POM—a Dimensionless Ratio)

The SI-POM tells CRE/facilities design engineers, technicians, and managers how much of a data center’s site power is consumed in overhead instead of making it to the critical IT equipment. This ratio is defined as:

$$SI-POM = \frac{\text{Data center power consumption at utility meter}}{\text{Total hardware AC power consumption at the plug for all IT equipment}}$$

SI-POM thus captures all the conversion losses in the transformers, UPS, PDU, and in critical power distribution losses, as well as cooling systems²⁶, lights, plus other minor building loads²⁷.

Determining SI-POM is important for two reasons.

- It can be reduced by improving the energy efficiency of the cooling, UPS and PDU system components. The building shell’s construction also plays some role. However, these opportunities occur either when a data center is first built or during major upgrading. Physical site infrastructure lasts for a long time, so component and system choices (good or bad) regarding SI-POM last a long time.
- The SI-POM value is very significantly influenced by operator skill in the running of the mechanical plant. *Institute* research indicates that focusing on improving skillful operation has a much bigger and more immediate impact than upgrading components once a data center has gone “live.”

It’s important to note that many design and operational features which save energy will not necessarily affect SI-POM. For instance, free cooling strategies using outside air (discussed later in this paper) are unlikely to improve SI-POM because a data center must still size cooling, UPS, and PDU systems large enough to handle peak cooling demands when free cooling is not available. The only way for free cooling to affect SI-POM is for free cooling to somehow (reliably) contribute to cooling year-round. This is unlikely, to the extent that peak cooling demands probably happen when a data center is running high compute load on a hot, humid summer day, while free cooling is mostly available on cool winter nights or in low humidity environments.

Example: Imagine a data center that draws a total of 1,000 kW at-the-plug AC IT hardware load. Due to conversion losses in the UPS and PDU systems, as well as the power consumed by the cooling system, 2,200 kW is consumed at the utility meter. The SI-POM is thus 2,200 / 1,000 = 2.2. This value is typical for a Tier IV data center.

²⁵ DC is a confusing community term because it can be used to mean “direct current” or “data center.” The *Institute* uses DC in the engineering sense to mean direct current and acknowledges that others use DC to mean data center.
²⁶ Here and elsewhere in the paper, the *Institute* uses the general terms mechanical or cooling systems to refer to all the equipment responsible for keeping data center temperature and humidity within the acceptable ranges. This includes the refrigeration consisting of compressors or chillers, cooling towers or dry coolers, pumps, cooling unit coils and fans, as well as equipment used for humidification and dehumidification and make-up air circulation.
²⁷ Note that if the data center uses direct current power distribution, the denominator for the equation would be hardware DC power consumption at the plug.

Energy: It's possible to create an analogous ratio using energy instead of power. The *Institute* calls this ratio the Site Infrastructure Energy Overhead Multiplier, or SI-EOM. In order to fully understand the efficiency of their site physical infrastructure, data center operators should track both SI-POM and SI-EOM separately.

Imagine, for example, a site which needs large chillers for summer cooling, but which also utilizes free cooling during winter months with little need for electrically driven chillers. This site would save a lot of energy during the winter, resulting in an SI-EOM for the season or the year that is significantly lower than its SI-POM during the summer. Free winter cooling may even become critically important in future data center site selection decisions.

Technical note: Since SI-POM is a ratio of power measurements, which change from instant to instant, SI-POM itself changes from instant to instant. There is a different value of SI-POM at each point in time. However, the main quantity of interest to data center users is the SI-POM measured at the instant of highest power draw. This will most likely be during a peak design day during the year, when both the IT equipment and the cooling systems are running at their highest loads. The issue becomes more complicated if for some reason the peak load of the IT equipment does not coincide with the cooling peak, for instance if computing loads are highest at night when it's cool outside. Luckily, since data center heat loads are dominated by internal equipment heat rather than outdoor conditions, it's unlikely that IT power demand would fail to coincide with peak cooling load in practice.

Strictly speaking, once installed, SI-POM will depend not only on operator skill in keeping the mechanical plant tuned-up, but also on where UPS and cooling equipment run in their operating curves. For example, a data center may have UPS, PDU, and mechanical systems that are very efficient at full load, but may be significantly less inefficient at low loads. If the data center has not yet been fully provisioned (which would be typical) and only runs at 30 percent of capacity, the resulting SI-POM may be quite high despite the apparent efficiency of the idealized equipment components²⁸. SI-POM also depends on ambient conditions, since it's partially determined by the mechanical plant efficiency, which can decrease as it gets hotter outside or as the outside relative humidity rises (depending on design configuration).

Note that while SI-POM changes instant to instant, SI-EOM only has a single value for the year, since it's a ratio of annual energy totals.

Finally, within SI-POM, the Institute believes it will be necessary to separately track power consumed by cooling and power lost in conversions within the UPS/PDU systems. Tracking these two quantities separately will allow data center users to more accurately estimate the at-the-meter savings from different types of infrastructure improvements.

IT Hardware Power Overhead Multiplier (H-POM, a Dimensionless Ratio)

H-POM tells IT procurement personnel how much of the power input to a piece of hardware is wasted in power supply conversion losses or diverted to internal fans, rather than making it to the useful computing components. For a single device, H-POM is the ratio:

$$\text{H-POM} = \frac{\text{AC Hardware Load at the plug}}{\text{DC Hardware Compute Load}}$$

For an entire data center, H-POM is the ratio of the total hardware load at the plug for the entire data center, to the total Hardware Compute Load for the entire data center²⁹.

Example: In the data center mentioned above with 1,000kW of Hardware Load at the plug, device power supplies are around 80 percent efficient, so about 20 percent of this power is lost in conversion. An additional five percent of the original power is consumed by fans within the hardware devices, so about 250kW is lost in total before reaching the pieces of the hardware that do useful computing. The H-POM is 1,000 / 750 = 1.33.

Energy: It's possible to create an analogous ratio using energy instead of power.

Technical note: As with SI-POM, H-POM is a power ratio measured at a particular point in time. There is a different value for H-POM at any given instant. Again, it's probably most appropriate for the measurement to be taken during the highest moment of compute load.

H-POM also depends on where the IT equipment runs on its operating curve. For example, a server may have a power supply that's very efficient at 100 percent of load, but inefficient at low loads. If the server runs only a single application that requires little electric power to the processor, the power supply could run very inefficiently, resulting in a high H-POM. Note that this scenario might be unlikely, because many servers that run at low compute loads still draw almost the same amount of power as when they run at full compute load.

DC Hardware Compute Load per Unit of Computing Work Done (Qualitative Only)

This metric helps IT architects and platform planners qualitatively describe how *power hungry* a particular platform is.³⁰ It excludes the differences in power supply and fan losses addressed by H-POM, instead focusing on how effectively a piece of equipment utilizes its internal DC watts. Some server architectures may consume more DC power than others for the same amount of computing work. Different spindle speeds of disk storage will consume different amounts of power per terabyte of storage provided; in many non-time-sensitive applications, the difference in latency won't matter.

²⁸ Rasmussen (2006) describes in detail how UPS system efficiencies can be especially low under low load conditions. This efficiency problem is compounded by Tier IV design topologies (See *Institute* white paper entitled *Tier Classifications Define Site Infrastructure Performance*, <http://www.uptimeinstitute.org/resources>) which under normal operation cannot have a power load exceeding 45 percent and typically operate between 15-30 percent loaded. This is one area where the business benefit of site infrastructure fault tolerance to assure information availability is in direct conflict with energy efficiency.

²⁹ Again, if the data center uses direct current power distribution, the numerator of the equation would be DC hardware load at the plug.

³⁰ Hennessy and Patterson (2007) describe the current dramatic shift in processing architecture from Instruction-Level Parallelism to Thread-Level Parallelism to exploit the efficiencies of multi-core processors for servers and workstations. Adding cores to increase performance may increase power consumption less than raising clock speed to increase performance.

For technical reasons, it is extremely difficult to create a single quantitative metric for this measurement. However, it’s important to clarify that vendors who wish to reduce the at-the-plug power use of their hardware can address this factor in addition to addressing H-POM.

Energy: It’s also possible to examine the energy (kWh) consumed per unit of computing work done.

Technical note: There are several major difficulties in creating a quantitative metric for DC Watts of Hardware Compute Load per unit of computing work. First, the term “computing work” means very different things depending on type of equipment: servers (processing data), storage (reading and writing data), or networking (routing packets). Second, even within a category like servers, what kind of computing work should be measured (web serving, databases, calculations, network throughput, etc.)? Third, there are numerous ways of quantifying computing work—processor cycles, MFLOPS, SPEC benchmarks, etc. For these reasons, the Institute describes this metric only qualitatively.

Deployed Hardware Utilization Ratio (DH-UR, a Dimensionless Ratio)

This metric helps IT executives quantify the fraction of their deployed IT equipment which is comatose—consuming energy/power while not running any application or handling data that anyone cares about. Unless specifically empowered to do so, most IT personnel will not risk turning devices off, since it’s often impossible to tell whether or not any given box is doing something important. Such dormant equipment can draw significant amounts of power as it idles 24 hours a day. For servers, the *Institute* defines the DH-UR as:

$$\text{DH-UR (Servers)} = \frac{\text{Number of servers running live applications}}{\text{Total number of servers actually deployed}}$$

The analogous definition for storage might look like the equation below:

$$\text{DH-UR (Storage)} = \frac{\text{Number of terabytes of storage holding important, frequently accessed data (within the last 90 days)}}{\text{Total terabytes of storage actually deployed}}$$

Clearly, data center management may not know these ratios exactly. (If comatose equipment could be easily identified, it would have already been turned off and removed.) Still, managers can often estimate these values or utilize industry average numbers when contemplating how much power and energy is wasted within their organization, at least as a starting point until the true data are available.

Example: Some data centers estimate that 30 percent of their servers are not doing anything productive. The live equipment ratio for servers in such an organization would be $(1-0.30) / 1 = 0.70$. If servers use 60 percent of the power in the data center, with the balance consumed by storage or networking equipment, this means that about 18 percent of the expensive UPS and cooling capacity in the data center is being wasted. A program to identify and shut off this load as a way of buying back data center capacity is far cheaper than building a new facility.

Technical note: In order to create a ratio that translates directly into power use, the best definition of DH-UR for servers would be:

$$[\text{Power at the plug consumed by live servers running important applications}] / [\text{Power at the plug consumed by all servers deployed}]$$

This definition would take into account the fact that idle comatose servers may draw less power than servers running applications, plus the fact that different models of servers use different amounts of power.

However, as noted earlier, idle servers often draw nearly as much power as servers doing real work, and it may be reasonable to simplify by assuming some level of average power consumption common to all models of server.

Under these two simplifying assumptions, the power-based definition above is equivalent to the simple ratio of numbers of servers.

A similar argument can be made about simplifying the DH-UR for storage in terms of terabytes.

Deployed Hardware Utilization Efficiency (DH-UE, a Dimensionless Ratio)

DH-UE helps IT managers quantify the opportunity for servers and storage to increase their utilization by virtualizing. For example, to ensure stability, many servers run only a single application, with the result that they operate at only 25 percent or less of their maximum compute load. But since servers running at low compute loads often draw nearly as much power as those running at high loads, a large number of partly-loaded servers quickly consumes valuable UPS and HVAC capacity (and raises the electricity bill).

For servers, DH-UE is defined as:

$$\text{DH-UE (Servers)} = \frac{\text{Minimum number of servers necessary to handle peak compute load}}{\text{Total number of servers deployed}}$$

The minimum number of servers necessary to handle peak compute load is defined as: the sum of the highest percentage of compute load each server experiences, plus any overhead incurred in virtualization, expressed as a percentage of the maximum compute load of a single server. Note that this calculation uses peak percent of compute load, not average percent compute load. A similar definition can be developed for storage.

Example 1: Imagine that a data center has nine identical servers which each run at 10 percent of compute load most of the time, increasing to 20 percent of compute load during peak business hours. Imagine also that the virtualization software is such that five percent of a server’s compute load is consumed in overhead for each virtual server that runs on the box. If the data center virtualizes all nine servers, the minimum number of boxes needed is: $9 \times .20 + 9 \times .05 = 2.25$.

So in reality, the data center would need at least three physical boxes running. Before virtualization, the DH-UE is: $2.25 / 9 = 0.25$.

After virtualization, the DH-UE is: $2.25 / 3 = 0.75$.

Example 2: It is critical to remember that data centers exist to serve a business purpose and business value is typically highest during periods of peak demand for services. To use average compute load in this calculation can have very negative performance impacts and therefore business consequences. IT managers must apply judgment on the highest percent compute load a server experiences,³¹ rather than the average percent compute load at which the server operates, when estimating virtualization potential.

Imagine that a data center has five identical servers, each of which runs at 15 percent load half the time (off-hours) and 25 percent load the other half of the time (business hours). The average percentage of compute load of each server is thus 20 percent. Assuming no virtualization overhead, it is only possible to virtualize four of the servers onto a single box, since that box would be running at 100 percent compute load during business hours. Despite the 20 percent average utilization, the data center can’t virtualize all five servers onto a single box, because it would run at 75 percent load during off-hours and be overloaded at 125 percent load during business hours. If a server becomes overloaded in this way, customers may experience unacceptable delays in the services they expect.

Technical note: The Institute recognizes that a precise definition of the percentage of compute utilization may not be obvious, and the authors invite contribution from both vendors and users on this point.

It’s also important to note that in many cases compute performance can fall off dramatically as compute load rises above some utilization threshold, often as low as 50-60 percent. One possible reason for this is that a limiting resource like RAM may become exhausted long before the processor is taxed to its full computing ability. So, any consideration of the “maximum potential for virtualization” must be sensitive to not just what is possible, but what is possible given the performance requirements dictated by business needs. The Institute looks forward to feedback from industry stakeholders to help refine this concept.

In order to create a ratio that translates directly into power use, the best definition of DH-UE for servers would be:

[Power at the plug consumed by the minimum number of servers necessary to handle peak compute load] / [Power at the plug consumed by all servers deployed]

This definition would take into account the fact that servers running at low compute load may draw less power than servers running at full compute load, and the fact that different models of servers use different amounts of power.

However, as noted earlier, un-virtualized servers running at low compute loads often draw almost as much power as virtualized servers running at full compute load, and it may also be reasonable to assume some level of “average” power consumption common to all models of server. Under these two simplifying assumptions, the power-based definition above is equivalent to the simple ratio of numbers of servers.

Free Cooling (kWh)

The free cooling metric allows facility designers and CRE/facilities decision makers to estimate the amount of money³² that could be saved each year by using cold outside air to directly or indirectly cool computer rooms, rather than using energy to run compressor-based refrigerant cooling systems.³³ Depending on geographic location, such free cooling can seasonally save a significant amount of energy (kWh) at the utility meter.³⁴ However, unless the outside air has the correct year-round wet bulb temperature range to meet a data center’s entire cooling load (no populated place in the continental US meets these requirements), site investment in compressor or chiller based cooling will still be required.

The free cooling potential is defined as:

$$\text{kWh saved each year at the utility meter} = \frac{\text{thermal kWh of free cooling used per year} \times \text{kW consumed at utility meter per kW delivered to cooling system}}{\text{COP of mechanical plant/cooling system}}$$

³¹ Compute performance and response times can fall off dramatically above some level of utilization which varies by platform. Determination of this inflection point is part of the “black magic” of IT capacity planners.

³² And also carbon, if the organization has a carbon reduction target.

³³ There are also many innovative types of free cooling such as the concept of storing-up enormous piles of ice/snow/water slush during the winter and using this thermal storage instead of electrical refrigeration for cooling in the summer

³⁴ The availability of free cooling may one day influence the choice of geographical location for data centers, although free cooling potential must be balanced against skilled labor availability and any extra telecom latency in distant locations.

The Coefficient of Performance (COP) of the mechanical plant/cooling system is defined as the thermal kW of cooling which can be delivered per electrical kW of input power. Note also that the kWh of free cooling used per year refers to *thermal* kWh, not electrical kWh.³⁵

There are some important technical limitations regarding free cooling:

1. It can only be used in certain climate zones, and only during hours of the year when outside temperature and humidity conditions are appropriate. Free cooling can be used in regions where humidity is low (Southwest). It can also be used during winter seasons when outdoor temperatures are low.
2. Free cooling does not save electrical or mechanical plant capacity, only energy. This is because the cooling system (and the power infrastructure to support it) must still be sized large enough to handle peak cooling loads at times when the use of free cooling is not feasible.
3. The direct use of outside air is not advisable in a Tier IV facility, where absolute reliability must be maintained even in the event of extremely unlikely outdoor disasters or accidents.

A freak dust storm, accidental chemical release, etc. could introduce contaminants and pollutants into the computer room from the outside air that could directly cause an outage or intermittent ghosts and perturbing soft errors. Should this happen, hardware manufacturers would likely not stand behind their fixed maintenance contracts. For Tier I and II facilities, however, the cost savings from direct outside air might make it worthwhile to bear the small risk of these extremely improbable events. Even for Tier IV facilities, there are mechanical system configurations that allow the indirect use of outside air for cooling without actually introducing it into the computer room.

Why is free cooling explicitly included in the *Institute's* metrics, rather than simply being lumped into the overall SI-EOM? When examining the key sub-factors affecting the SI-EOM, free cooling is one of the biggest variables depending upon certain climates. Ignoring free cooling as a major element in SI-EOM metric benchmarking risks placing data from all data center facilities in the same bucket, even though sites located in areas with opportunities for significant free cooling should be held to a higher standard. (If all sites were lumped together, a site with an otherwise very inefficient infrastructure could still have a misleadingly low SI-EOM if it used a lot of free cooling.) So, anyone assessing data center efficiency must consider the use of free cooling wherever appropriate.

Example: Imagine a data center with 1,000kW of deployed IT equipment, which generates 1,000kW of heat. Assume also that for around a quarter of the hours in a year the facility can meet 50 percent of its cooling needs with outside air, with refrigerant

compressor cooling picking up the rest. The mechanical system has a COP of 2.5 indicating that it can remove 2.5kW of heat from the building for each kW of input electricity. The chilled water storage portion of the mechanical system is also a critical load backed by UPS with some applicable conversion losses, so 1.05kWh are consumed at the meter for each 1kW that makes it to the mechanical system.

Overall, then, free cooling supplies: $1,000 \times 8,760 \times 0.25 \times 0.50 = 1,095,000$ kWh of cooling.

And the number of kWh saved, per year, at the utility meter is: $1,095,000 / 2.5 \times 1.05 = 459,990$ kWh.

Technical note: Mechanical engineers often think in terms of BTUs for cooling, but converting BTUs to kWh makes the definition of the metric simpler. Operators might think of cooling in terms of kWh for cooling anyway, since they know IT load kW (which all becomes heat) and they know 8,760 hours/year.

COP is a common measure of cooling system efficiency, defined as kW of cooling delivered per kW of input power required. Pieces of equipment may instead list efficiency as kW/ton of cooling, EER, SEER, etc., but all of these other metrics can be converted to COP. Note that COP also fits the definition of efficiency (power out / power in), but COP can exceed 1.0 for heat pumping systems. The Institute doesn't wish to create confusion by referring to "efficiencies" greater than 100 percent, so the free cooling metric uses COP instead.

Technically, the COP of a data center mechanical system will vary from instant to instant depending on the load. Also, depending upon design configuration, during times when free cooling can only deliver part of the needed cooling, the chiller must still operate. In these cases, the chiller could be operating at a very low load, which reduces its efficiency, offsetting some of the energy savings from the free cooling. The Institute neglects to account for both these complexities here.

Strictly speaking, the final term in the equation should be:

[kWh consumed at the utility meter per kWh delivered to the mechanical system].

Using kW is an acceptable proxy, although technically they are not the same since the ratio of kW at the meter to kW delivered to the mechanical system can vary, due to differing efficiency losses at different loads.

Note also that this final term is not the same as SI-POM, because data center mechanical equipment has fewer layers of conversion losses between itself and the meter than does IT equipment. (i.e., the mechanical plant may or may not be behind a UPS.)

³⁵ One kWh of thermal energy is equal to 3,412 British Thermal Units (BTU).

Enable Energy Saving Features (kWh)

The energy-save metric allows IT analysts to estimate the amount of energy, money, and carbon that could be saved each year by letting IT equipment hibernate during times when it is not being used. The number of kWh which can be saved by hibernating idle equipment is:

$$\text{kWh saved each year at the utility meter} = \left[\begin{array}{cc} \text{kW drawn by one} & \text{kW drawn by one} \\ \text{piece of equipment -} & \text{piece of equipment} \\ \text{when idle} & \text{when hibernating} \end{array} \right] \times \begin{array}{c} \text{Number of pieces} \\ \text{of equipment} \\ \text{that could hibernate} \end{array} \times \begin{array}{c} \text{Number of hours} \\ \text{per year equipment} \\ \text{could hibernate} \end{array} \times \text{SI-POM}$$

It is important to note that energy-save features do not save electrical power capacity or mechanical plant capacity, only energy. This is because site power and cooling systems must still be sized to handle the times of peak compute demand.

Also, many data center operators have expressed frustration with energy-save features that do not work correctly. Often, IT equipment fails to power-down when it should, or (even worse) fails to power up when it should.³⁶ In order to ensure system stability, many IT shops disable energy-save features as one of the first steps to deploying a new piece of equipment. However, the Institute’s research indicates that the savings which could be realized from effective and reliable energy-save features could be substantial in many cases, and therefore encourages data center decision makers to require improvements to these features from vendors and to mandate their use within their IT organizations. Simply abandoning the idea of energy-save features leaves a potentially large efficiency opportunity untapped.

Example: Imagine a data center with 2,000 servers, each of which draws 0.5kW when running at full compute load, for a total of 1,000kW of IT equipment installed. The facility is Tier IV and has a SI-POM of 2.2. Assume that 800 of the servers are needed all 8,760 hours of the year, but that the other 1,200 are idle for 4,576 hours per year. (This roughly corresponds to being idle during the eight night-time hours of weekdays and all day on weekends.) When idling, these 1200 servers would draw 0.30kW, but in hibernate mode they draw only 0.05kW. The amount of power saved at the utility meter each year by allowing these servers to hibernate is $(0.30 - 0.05) \times 1,200 \times 4,576 \times 2.2 = 3,020,160$ kWh.

Technical note: The Institute uses the term “hibernate” generically here and is aware that there are specific industry-accepted definitions of hibernate, sleep, power-save mode, etc., all of which fall in a different spot on the tradeoff spectrum between saving power and processing latency on quickly returning to life. Also, it’s not clear that these definitions are consistent between different types or even brands of equipment.

Many pieces of hardware can also enter multiple energy saving modes. There may be times when operators enable a “light sleep” mode to save energy during some periods and a “deep sleep” mode for further savings during periods when a slower resume of service is acceptable.

To make this metric usable, the Institute has made a number of simplifications here. It’s clear that all models and types of IT hardware are not identical, nor will the difference between their idle and hibernate power consumption be identical. Different pieces of equipment may not even be hibernating at the same time. What one really wants to know is the total kWh saved if each piece of equipment was hibernating every moment that it could. This would be:

$$\sum_{\text{all pieces of equipment } i} \left[\int_{\text{all possible hibernate time}} (\text{idle power}_i - \text{hibernate power}_i) (\text{SI-POM}) dt \right]$$

However, it’s much easier to just assume some sort of average power savings per piece of IT equipment, the number of pieces of equipment, and the number of hours idle.

This is further complicated by the fact that SI-POM will vary from hour to hour, since efficiencies of the UPS, PDU, and mechanical systems will vary with load. The mechanical/cooling systems will also operate more or less efficiently depending on ambient conditions. The Institute ignores these complexities as well.

Technically, it is possible that hibernating IT equipment could save power and cooling capacity, but only if two sets of equipment are needed at different times, (i.e., one runs while the other hibernates). It would have to be the case that both sets of equipment were never needed at the same time.

³⁶ Virtualization of servers, in which multiple instances of an operating system are active at different times, only complicates matters further.

Putting it All Together

The next high-priority area for research and discussion is to lay out explicitly how these metrics nest together, so that data center users can estimate the effect of improving various efficiency metrics on their overall at-the-utility-meter electricity consumption (both peak power and total energy). This will allow user executives and public policy makers to prioritize sub-areas for improvement and to more easily determine benchmarks for “good” versus “below average” levels of efficiency. The *Institute* looks forward to the discussion of nesting methodology at the *Institute’s* Design Charrette on Data Center Efficiency (October 28-30 in Santa Fe, NM) and afterward.

The *Institute* looks forward to continuing dialogue with all stakeholders in the data center space regarding these metrics and the criteria proposed to evaluate them. The authors hope this paper represents a significant and useful addition to the ongoing discussion.

Conclusion and Next Steps

The *Institute* hopes that this paper has provided decision makers with a sufficiently senior-level view of efficiency metrics to enable rigorous thinking about how to obtain the most business value from IT hardware and data center investments, and about how optimizing these metrics can contribute to organizational environmental goals. This paper is intended to build upon previous metrics work and to move data center operators, equipment manufacturers, and corporate and public policy makers toward a general consensus regarding energy/power efficiency metrics and how a green data center should be defined.

In particular, the authors intend for readers to take away the following points:

- Any discussion of metrics should include some criteria for evaluating them, including intuitive use, technical accuracy, flexibility in the face of new developments, and vendor neutrality.
- Metrics will work best when they divide data center efficiency into manageable chunks that can be clearly assigned to different groups within an organization. To be most effective, however, these different groups must work together closely rather than optimizing in isolation.
- Power and energy should be considered separately, both for technical accuracy and to clearly define whether efficiency measures yield benefits for data center capacity, actual energy reduction, CapEx and OpEx reduction, carbon footprint reduction, or all of the above.
- Data center decision makers can most effectively address efficiency by considering what configurations meet their present and projected business needs, rather than merely looking at how to make existing configurations slightly more efficient.
- Even metrics which seem simple at first glance contain hidden assumptions and simplifications. It is important to list these explicitly so that they are not forgotten.
- Even after consensus is reached on efficiency metrics at a meta level, there is still considerable (and possibly contentious) work to be done around the exact technical and tactical procedures for measuring them.

References

- Atkinson, Robert D., and Andrew S. McKay. Mar 2007. *Digital Prosperity: Understanding the Economic Benefits of the Information Technology Revolution*. Washington, DC: The Information Technology and Innovation Foundation. Available at <http://www.itif.org/index.php?id=34>.
- Brill, Kenneth G. Aug 2007. *The Invisible Crisis in the Data Center: The Economic Meltdown of Moore's Law*. Uptime Institute, Santa Fe, NM.
- Brill, Kenneth G. Jan 2007. *Data Center Energy Efficiency and Productivity*. Uptime Institute, Santa Fe, NM.
- Brill, Kenneth G. Mar 2006. *High Density Computing: The Path Forward 2006*. Uptime Institute, Santa Fe, NM.
- Calwell, Chris. 10 July 2007. *Personal communication with John Stanley of Uptime Institute*. Chris Calwell is VP of policy and research at Ecos Consulting, and has multiple publications in the field of electronics power consumption.
- Green Grid. 20 Feb 2007. *Green Grid Metrics: Describing Datacenter Power Efficiency*. Available at http://www.thegreengrid.org/gg_content/Green_Grid_Metrics_WP.pdf.
- Hennessy, John L. and Patterson, David A. 2007. *Computer Architecture: A Quantitative Approach 4th Edition*. Morgan & Kaufmann.
- Koomey *et al.* October 2007. *A Simple Model for Determining True Total Cost of Ownership for Data Centers*. Uptime Institute, Santa Fe, NM
- Koomey, Jonathan. 2007. *Estimating total power consumption by servers in the U.S. and the world*. Oakland, CA: Analytics Press. February 15. Available at <http://enterprise.amd.com/usen/AMD-Business/Technology-Home/Power-Management.aspx>.
- Koomey *et al.* 3 Nov 2006. *Server Energy Measurement Protocol*. Oakland, CA: Analytics Press. Available at <http://www.energystar.gov/ia/products/downloads/Finalserverenergyprotocol-v1.pdf>.
- Malone, Christopher, and Christian Belady. 19 Sep 2006. *Metrics to Characterize Data Center and IT Equipment Energy Use*. Proceedings of 2006 Digital Power Forum, Richardson, TX, September 19, 2006.
- Popik, Drew. 12 July 2007. *Email communications with John Stanley of Uptime Institute*. Drew Popik is president of Data Structure Management Corporation.
- Rasmussen, Neil. 2006. *Electrical Efficiency Modeling of Data Centers*. APC. White paper #113. Available at http://www.apcmedia.com/salestools/NRAN-66CK3D_R1_EN.pdf
- SPEC. 2007. SPEC Power and Performance Committee. Standard Performance Evaluation Corporation. Available at <http://www.spec.org/specpower/>
- U.S. Environmental Protection Agency, ENERGY STAR Program. 2 Aug 2007. *Report to Congress on Server and Data Center Energy Efficiency* (Public Law 109-431). Available at http://www.energystar.gov/index.cfm?c=prod_development.server_efficiency.

White Paper Reviewers

Each subsequent draft of this paper is being publicly reviewed and will continue to evolve as an open-source document. We appreciate all feedback. Many constructive comments came in too late to be incorporated into this release, but all reviewers will receive a response. Among the reviewers providing comments on this white paper were Dave Perkins, Active Power, Inc.; Stan Burchell, Abbot Laboratories; Guy Majestic, Barclays Capital; Jeff Trower, Data Aire, Inc.; David Moss, Dell; Patrick Donnelly, R. Patrick Donnelly and Associates; Kathrin Winkler, EMC Corporation; Dave Leonard, Infocrossing; Lex Coors, InterXion; Thomas Delaney, JPMorgan Chase; Rick Salmon, Lehman Brothers; John Sasser, Sabey Corporation; John Dennis, Sutter Health; and Michael Siegele, VISA.

About the Authors

John R. Stanley was the *Institute's* 2007 Summer Intern working under the supervision of Kenneth Brill, Executive Director. He is currently an M.S./Ph.D. candidate in the Energy and Resources Group at the University of California, Berkeley. During spring 2007, he worked on the research team at Lawrence Berkeley National Laboratory which assisted the EPA in preparing their report to Congress on data center efficiency. Prior to graduate school, Stanley worked at Rocky Mountain Institute, where he served on consulting teams analyzing energy efficiency opportunities for a variety of private sector commercial and industrial clients.

Kenneth G. Brill is the founder and Executive Director of the *Institute* and its 100-corporate member Site Uptime Network in both North America and Europe. He holds an undergraduate degree in electrical engineering and an MBA from the Harvard Business School. Many industry innovations including dual power and the industry's Tier system for evaluating data center design level trace back to his original conceptual work. In 1999, recognizing that increasing heat density would become critical to IT availability; Mr. Brill worked closely with the Thermal Management Consortium to publish the first *2000-2010 Heat Density Product Trends* white paper. This foundational document predicting many of the problems now facing IT executives was updated in April 2006. He has authored or contributed to many white papers and is a frequent commentator in business and technology media on data center and site infrastructure design, engineering, and management issues. His current focus is on strategic and business impacts of The Economic Meltdown of Moore's Law.

Jonathan Koomey, PhD is a Project Scientist at Lawrence Berkeley National Laboratory and a Consulting Professor at Stanford University, as well as an *Institute* Senior Fellow. Dr. Koomey is one of the leading international experts on electricity used by computers, office equipment, and data centers. His study, *Estimating total power consumption by servers in the U.S. and the world*, published in February, 2007 is already a standard reference work for anyone working in the field of data center energy. He has also published extensively on critical thinking skills. He holds M.S. and Ph.D. degrees from the Energy and Resources Group at the University of California at Berkeley, and an A.B. in History of Science from Harvard University. In 1993 he won the Fred Burgraff Award for Excellence in Transportation Research from the National Research Council's Transportation Research Board. He was named an Aldo Leopold Leadership Fellow in 2004 and an AT&T Industrial Ecology Fellow in January 2005.

About the Uptime Institute, Inc.®

Since 1993, the Uptime Institute (*Institute*) has been a respected provider of educational and consulting services for Facilities and Information Technology organizations interested in maximizing data center uptime. The *Institute* has pioneered numerous industry innovations, such as the Tier Classification System for data center availability, which serve as industry standards today. The *Institute's* 100 global members of the Site Uptime Network® represent mostly Fortune 100 companies for whom site infrastructure availability is a serious concern. They collectively and interactively learn from one another as well as from *Institute*-facilitated conferences, site tours, benchmarking, best practices, and abnormal incident collection and analysis. For the industry as a whole, the *Institute* publishes white papers, offers a Site Uptime Seminar Series and a Site Uptime Symposium and Data Center Design Charrette Series on critical uptime-related topics. The *Institute* also conducts sponsored research and product certifications for industry manufacturers. For users, the *Institute* certifies data center Tier level and site resiliency.



Building 100, 2904 Rodeo Park Drive East, Santa Fe, NM 87505
505.986.3900 · Fax 505.982.8484 · uptimeinstitute.org