UNIVERSITY OF OREGON
APPLIED INFORMATION MANAGEMENT

Presented to the Interdisciplinary Studies Program:
Applied Information Management and the Graduate School of the University of Oregon
in partial fulfillment of the requirement for the degree of Master of Science

Lowering the Carbon Emissions Footprint of Enterprise Data Centers Through Energy Efficiency Gains

CAPSTONE REPORT

Cecilia Tomory
Principle DBA
PacifiCorp

University of Oregon
Applied Information Management
Program

February 2010

Continuing Education
1277 University of Oregon
Eugene, OR 97403-1277
(800) 824-2714

Approved by

Dr. Linda F. Ettinger
Academic Director, AIM Program

RUNNING HEAD: Lowering DC Carbon Emission

Lowering the Carbon Emissions Footprint of Enterprise Data Centers Through Energy Efficiency Gains

Cecilia Tomory

PacifiCorp

This page intentionally left blank

Abstract

Selected literature published between 2004 and 2009 is mined for emerging trends and best practices regarding green IT and data centers. Environmental issues are becoming a serious industry concern, spurred by local and global initiatives (Kumar, 2009). Steps are identified for enterprise data center managers, in order to measure, manage, and improve power consumption efficiency within their facility. Data center energy efficiency gains save power, lower energy costs, and decrease overall corporate carbon footprint (Freeman, 2009).

This page intentionally left blank

Table of Contents

This page intentionally left blank

List of Figures

This page intentionally left blank

Introduction

Problem Area

According to the United States Department of Energy (USDOE) computer-based models have shown that an increase in greenhouse gas emissions results in the gradual rise of the Earth's average surface temperature. It expects that the rising temperatures over time may produce a climate change--that is, changes in weather patterns, sea levels, and storm severity. The USDOE also states that 82 percent of the total carbon dioxide greenhouse gas emissions in 2006 were generated by the combustion of petroleum, coal, and natural gas (2008).

Gartner, Inc. is an internationally recognized research analysis company specializing in the IT industry (Routers, 2009). In 2007, Gartner found that information communications and technology is responsible for approximately two percent of global carbon dioxide (CO2) emissions, which is as much as the aviation industry produces. Gartner also believes this rate of emission is unsustainable. Furthermore, Gartner, Inc. followed up this first statement later in the year declaring that the energy consumed by data centers is responsible for almost a quarter of global carbon dioxide emissions produced by information and communications technology (October 2007).

Purpose

The purpose of this study is to identify then compile emerging trends and best practices into a set of practical recommended modifications that when applied within enterprise data centers may achieve lower carbon emissions. The focus of the study is on specific changes that are practically implementable within existing facilities, rather than the management of the change process. Because data centers are facilities that host hardware which process information, facilities are categorized by size and availability goals rather than the type of information, or industry, they support. Therefore, this study does not categorize results by industry but rather by data center tier where it might be of relevance.

The study is designed as a literature review (Leedy & Ormrod, 2005). Literature published between 2004 and 2010 is examined in the following content areas: (a) measuring data center power efficiency (Mathew, Greenberg, Ganguly, Sartor & Tschudi, 2009), (b) data center design and technology energy efficiency (Snevely, 2002), and (c) carbon footprint or emissions (Wiedmann & Minx, 2007) as related to the Information Technology industry.

An enterprise data center is defined as a Tier III or Tier IV data center. This tiered rating scale was developed by the Uptime Institute and defined by Turner, Seader and Brill (2005) as: A Tier III facility is "composed of multiple active power and cooling distribution paths, but only one path active, has redundant components, and is concurrently maintainable, providing 99.982% availability" (p. 2) and a Tier IV facility is "composed of multiple active power and cooling distribution paths, has redundant components, and is fault tolerant, providing 99.995% availability" (p. 2).

Significance

The sole purpose of a data center is to support information technology and communications (USEPA, 2007), which in turn supports information management. The effort to decrease energy consumption and the overall carbon footprint of data centers should be a goal for the industry as a whole (Baroudi, Hill, Reinhold, & Senxian, 2009). As the issue of global warming continues to escalate in priority for various national governments, the chance of new legislation requiring emission limits on data centers is becoming a distinct possibility (Ebbers, Galea, Schaefer & Khiem, 2009). According to The Green Grid:

> Climate change and fuel security have dramatically risen up the policy agenda, resulting in a comprehensive policy framework at both the European Union (EU) and national levels. The growth of data centers, and their energy intensity, means the industry will be particularly affected by the drive for energy efficiency of product design and building services. (2009, p. 1)

According to the Hewlett Packard Company (HP) there is also indication that policy changes in the United States directly impacting data center energy utilization may be forthcoming. HP states that federal public law passed December 2006 requiring the United States Environmental Protection Agency (USEPA) to study, promote, and recommend incentives to advance energy efficiency in datacenters (2009).

The USEPA claims almost all economy sectors from financial and academic, to government institutions employ data centers (2007). The underlying assumption of this study is that a very small percentage of these businesses will be willing to significantly rebuild their data centers to incorporate a broad spectrum of new technologies and processes. The latest IT spending prospective published by Gartner, Inc reinforces the lack of appetite for corporate IT spending in the near future:

> Hardware projects continue to be stalled for PCs, servers and storage, further pushing down the new sales of infrastructure software that are dragged by hardware sales. Also, new sales of enterprise application software in the manufacturing and financial sectors

> have completely stalled as these vertical sectors sort out their long-term viability. (2009, p. 7)

Many organizations will be looking for small incremental changes to help lower emission cost without a large capital investment and without disruption to their daily business. In effort to minimize expenditures, data centers are attempting to postpone replacement of hardware and facility infrastructure components (Nexsan, 2009).

Audience/Outcome

This study is designed for the information management executives in these organizations who are directly responsible for the capital budget investment in their data centers, as well as the ongoing overhead costs of operating the data centers. A secondary audience is any information management professional who is interested in influencing their corporations in adopting emerging data center energy efficiency practices.

The review highlights the latest research addressing the energy consumption problem and emerging practices available for existing tier III and tier IV data centers. The intent is to produce a concise summary of industry trends and recommended best practices developed within the past five years, in each element of this study: green IT and carbon emissions, measuring data center energy consumption, data center design and technology improvement for increased efficiency of energy consumption. Recommended modifications which can be made within enterprise data centers, along with best practices, are grouped by study element. The sub-grouping of recommendations allows for addressing each element individually. Segmenting the larger task of lowering data center carbon emission through energy efficiency gains, in three distinct elements, may precipitate a more palatable gradual implementation of changes. The summary of emerging trends, best practices, and recommendations provides a comprehensive overview for the target audience, a starting point to perhaps a longer journey of change.

Delimitations

The following delimitations frame the study:

- Each item in the selected literature addresses at least one element of the study: energy efficiency of data centers, data center infrastructure and design, measuring and lowering resource consumption of a data center, green IT.
- All sources cited in the literature can be located and/or referenced directly.
- The literature meets minimum selection and evaluation criteria based on the authority of the author and publisher, objectivity of the author, the quality of work, the coverage of work and currency. Each of these areas is evaluated as outlined by the University of Oregon Library's critical evaluation of information sources (Bell & Smith, 2009).
- If the author or publisher of the literature does not meet the standard evaluation criteria as outlined by Bell and Smith, it is included if it can be determined that the author is an experienced professional in the field of either: data center design, data center technology or infrastructure energy consumption, the measurement or benchmarking of energy consumption, server virtualization, sustainability, carbon emissions.
- The literature included in this study is published no earlier than January 1, 2004. There is very limited relevant literature available from any source meeting the minimum criteria, as outlined above, prior to this date. According to the USEPA, the exponential increase in energy consumption of data center servers began in 2002 (Aug, 2007).

- Literature on the topic of *carbon footprints* uses the term to indicate a measurement of greenhouse gasses expressed as carbon dioxide equivalent caused by an activity or accumulated over the life of a product (Wiedmann & Minx. 2007).

Data Analysis Plan Preview

The literature selected for review is divided first by major topic and then by source. The four source categories are academic papers, government publications, vendor papers, and journal pieces. Each source category is then analyzed in relation to the three major topic categories: green IT and carbon emissions, measuring data center energy efficiency, data center design and technology energy efficiency. Ideas presented by each of the four sources of writings, across each of the categories, are analyzed in accordance with the conceptual analysis process described by Busch, De Maret, Flynn, Kellum, Le, and Meyers (2005). This process provides a research tool which can be used to verify the presence of certain words or concepts within texts or sets of texts. The overall data analysis goal is to identify best practices and recommended modifications that will help an established enterprise data center lower its carbon footprint by achieving higher overall energy efficiencies.

Data handling and documentation. Selected literature is scanned for possible relevance, grouped into categories, and prioritized material is read to determine order of importance (Obenzinger, 2005). Two primary tools are used for recording priority, key ideas, and useful quotes. An Excel spreadsheet tracks the literature collected for the study. For each piece the table includes its reference, the source of the material, its relevance to each of the three elements of the study, and the final priority of the piece. A word document is used to make notations along with copies of useful quotes from each of the selected literature. Notations for each material include: sub questions addressed, how key terms are used, major trends, any statistics contained in the literature, emphasis, strengths and weaknesses, any gaps, key descriptors (Mongan-Rallis, 2006). An Entity Relationship diagramming tool is used to map identified key concepts and trends, to identify possible relationships.

Writing Plan Preview

The Review of the Literature section of the paper is structured by a thematic rhetorical strategy, organizing research around ideas or insights (Literature Reviews, n.d.). The review is composed of three themes, each relating to one of the three key subtopics of the study. Information is presented in a manner that reveals the development of each topic over time. The goal is to show how the field of data center management and sustainability is developing and to report the emerging best practices for reducing data center carbon emissions through energy usage efficiency gains. The Conclusion of the review, based on further analysis of the themes and emerging trends, identifies best practices along with modifications most commonly agreed upon in the literature, which an enterprise data center may implement to lower the carbon footprint through by achieving higher overall energy efficiencies.

Definitions

This study reviews a selection of published literature relating to a set of technical issues including the measurement, management, and reduction of energy consumption within a data center. The terminology used in the selected literature, as well as throughout this study, has specific meaning in the context of information technology and data centers. This section of the study defines words, phrases, and acronyms by stating their meaning within the scope of the study (Definitions, n.d.). The definitions provide a common vocabulary.

Best practices. Processes and methodologies that have consistently shown better results than those achieved by other means, and are used in the industry as results to strive for (Best practices, n.d.).

Carbon emissions. The release of carbon dioxide (CO2) into the earth’s atmosphere (United States Government, 2008)

Carbon footprint. The measurement of greenhouse gasses emitted by a specific entity over a stated period of time and expressed as units of carbon dioxide equivalent (Carbon footprint, n.d.; Carbon Footprint Ltd, 2009).

Chiller. A subcomponent of air conditioning systems for large buildings. Chillers produce cold water to remove heat from the air in the building typically by mechanically compressing using a refrigerant (Chiller, n.d.).

Computer server. as defined by the USEPA Energy Star program: “A computer that provides services and manages networked resources for client devices, e.g., desktop computers, notebook computers, thin clients, wireless devices, PDAs, IP telephones, other Computer Servers and other networked devices" (April 2009, p. 5).

Computational fluid dynamics (CFD). Models the airflow, temperature, and pressures at a room-level. It is used to analyze overall impact of hardware modifications on the thermal environment within the room (Long, Freeman & Patterson, 2009).

CRAC . Computer room air conditioning units. "A device that monitors and maintains the temperature, air distribution and humidity in a network room or data center" (TechTarget, 2009).

Data center. A facility, or space, with power conditioning and cooling equipment, computer systems and associated components, that host production application systems (Data center, n.d.; USEPA, 2007; Avelar, Torell & Hampel, 2008).

Data center energy use. Is the total of all energy consumed by the facility including heating, cooling, and source energy (Mathew, Greenberg, Ganguly, Sartor &Tschudi, 2009).

Data center infrastructure. The site and facilities that support the data center including cooling, HVAC, UPS, power, standby generators, and alternative power sources (Ebbers, Galea, Schaefer & Khiem, 2009)

DCiE IT. Equipment power / total facility power . The closer the number is to 1 the greater the power usage efficiency within the data center (Ebbers, Galea, Schaefer & Khiem, 2009).

Deduplication. A storage based technology that optimizes storage resources by not storing redundant data. The storage of duplicate data is minimized through a combination of physical bit mapping, catalogs, and algorithms (Freeman, 2009).

Green Computing. The practice of applying computer resources in a manner that reduces toxic emissions while maximizing energy efficiency, while reducing and/or reusing any retired materials (Mainline, 2009).

HVAC. Heating, ventilation, and air conditioning including chillers, pumps and air handling units (USEPA, Sept, 2009).

Idle. The state a computer server is fully powered and operational, but not processing any useful work (USEPA, 2009).

Infrared thermography. Tool used to find hot spots and thermal patterns in the data center. It provides a temperature-based color picture of the surface temperatures in the room (Long, Freeman, & Patterson, 2009).

PUE. Power usage effectiveness. All power consumed by the data center / power consumed by IT equipment (Ebbers, Galea, Schaefer & Khiem, 2009).

Server consolidation Allows for multiple workloads to be processes on as few machines as possible (Ebbers, Galea, Schaefer & Khiem, 2009)**.**

Tier III data center. Systems are concurrently maintainable. All planned outages to these systems have no visible impact to the end users. Systems availability is no less than 99.982% (ADC Inc, 2006).

Tier IV Data center. Systems are fault tolerant. Both planned and unplanned outages to the systems have no visible impact to end users. Systems availability is no less than 99.995% availability (ADC Inc, 2006).

Trend Analysis. Analysis of changes over time (Princeton, 2009). More specifically, changes in methodology used to measure and monitor energy consumption within data centers.

UPS. Uninterruptable power supply unit. Referring to "a combination of converters, switches and energy storage means, for example batteries, constituting a power supply for maintaining continuity of load power in case of input power failure" (USEPA, 2009, p. 17).

Virtualization. "Eliminates the physical bonds that applications have to servers, storage, or networking equipment" (Ebbers, Galea, Schaefer & Khiem, 2009, p. 12).

Research Parameters

This section of the study outlines how the study is conducted. The research topic is driven by a problem statement or question, and sub questions. The research parameters define the search, documentation, analysis, and writing strategies employed in the study. Parameters limit the scope of study and guide the work such that the objectives are well understood and met. These parameters state what literature will be considered for inclusion and what will not. How the literature is identified, content is coded, and results are documented is included to ensure the consistent treatment of the literature.

Research Questions

Main question. What changes can enterprise data centers make to lower carbon emissions?

Sub-questions.

- What is a commonly accepted definition of a carbon footprint?
- What is green computing and what is a green data center?
- How is carbon emission metrics used as a measurement of a data center's carbon production?
- What physical facility changes can a data center make to lower carbon emissions?
- What technology changes can a data center make to lower carbon emissions?
- What data centers have been recognized nationally for making significant efforts to decrease carbon emissions in the past five years? Why?
- What have these recognized data centers implemented to gain the recognition?
- Is there any pending local and/or international government policy that may influence the need for energy efficiency gains within a data center?

Search Strategy Report

The research for this study focuses on three areas of literature: measuring data center emissions, data center design, and data center power consumption. Additional focus is added by intersecting these three topics with that of carbon footprints or emissions.

The initial search for literature is grouped in the following manner by types of sources: peer reviewed scientific research, government publications, vendor publications, and journalist pieces. The preliminary goal is to establish where the general thinking is on the issue and where the most applicable sources of information might be.

Search terms. Key words are derived from government literature, such as the USEPA report on Data Centers published in 2007, vendor white papers such as "The Green Data Center" published by IBM in 2009, and from IT industry standard terminology. For example the term Data Center is a term commonly used in the IT industry since the early 1990s (Data Center, n.d.).

The dominant method of search is to follow reference chains initiated in vendor white papers. Broad database searches using key words prove to be less effective since much of the relevant literature is very new and predominantly published by federal government agencies, vendors, or professional organizations. There is limited traditional academic literature published as of yet. The following key words are used:

- data center carbon footprint
- data center classification
- carbon emission
- green house effect
- DOE
- data center cooling
- infrastructure green
- green IT
- data center metrics
- data center power
- data center efficiency
- measuring energy efficiency
- data center thermal modeling
- monitoring energy consumption
- facilities design
- leed
- GEIT
- thermal workload scheduling

- server power
- data center energy consumption
- data center spending
- corporate IT spending
- data center budgets
- energy policy
- carbon policy

Search results. A broad range of people and institutions is trying to solve the data center carbon footprint dilemma. Available literature is produced by various government and academic institutions, vendor research labs, analyst groups, and "think tanks". Each is trying to tackle the issue from a number of angles, some by finding technology solutions to address power consumption and heat emissions of the hardware in question, others by proposing new ways to host information systems.

	Search Engine/Results			
	UO local	World Cat	Google Scholar	OneSearch Env Studies
carbon +	na	na	na	
Footprint	1070	1070	47300	7
Metrics	1	360	50700	1
Datacenter +	2250	20353	na	1
Emissions	1	1	16300	0
Carbon	0	20353	1520	0
Green	0	6	16700	0
metrics+carbon	0	0	3280	0
GEIT	617	617	19300	0
Uptime institute	37	37	23700	0
DOE + datacenter	5	1	16400	0
data center information technology costs			17,100	
data center cost OR budget OR 2010			35,000	
green information technology (IT) +	17	5971	11100000	0
Awards	0	60	82900	0

Table 1 - Search Terms & Results

Literature resources. Text, in the format of books, articles, and websites, are retrieved mainly via online search engines and databases. Initial search terms are derived through government literature on Sustainability, IT industry white papers, and journal articles. Searches are conducted using a number of search engines, with Google and Goggle Scholar producing the bulk of literature sources. UO Libraries Catalog and the IEEEE Computer Science Digital Library are the two primary databases mined for relevant content.

The most effective additional resource of literature is the reference lists from discovered literature. These bibliographies and references provide an extensive chain of interesting papers, articles, and web sites. Professional contacts throughout various vendor companies are also a great source of literature regarding specific vendor research and development. A number of professional organization websites are also sourced for key words and white papers and references, including the sites for The Green Grid (www.TheGreenGrid.org), Uptime Institute (www.UptimeInstitute.org), and the Energy Efficient Data Center Demonstration Project (dcee.svlg.org). The dominant method of search is following reference chains initiated in vendor white papers. Broad database searches using key words have proven very ineffective since much of the relevant literature is very new and predominantly published by federal government agencies, vendors, or professional organizations. There is very little traditional academic literature published as of yet.

Documentation Approach

Every selected literature is reviewed at a high level for relevance. If, based on the abstract, introduction, and perhaps references, it seems to be interesting then the reference and source of are recorded in a spreadsheet. An initial impression of material's relevance to the three primary elements of the study, measuring data center emissions, data center design, and data center power consumption, is also recorded. This spreadsheet is later used to track coding results. A word document is used to make notations and save interesting quotations. Finally, an Entity Relationship diagramming tool is used to map identified key concepts and trends, to identify possible relationships

Selection and Evaluation Criteria

Each selected reference is initially identified by a keyword search. The key words are those relevant to sustainability, carbon emissions, and energy consumption within a data center, as noted above. The searches are then limited by published date no older than January 1, 2004. This limitation on date is necessary to ensure that only current trends and best practices are included in the study. Once a reference is identified based on date and key words, the abstract and introduction is read to determine potential relevancy to one of the research papers key themes.

Each potentially relevant reference is then verified for credibility. If the document contains citations then a random sampling of the citations is checked. Any information given on the authors and/or publishers is examined. Any sources of funding for the paper and other possible sources of content bias are also identified. Research methods employed in producing the literature are analyzed.

In summary, all documents included in the review are relevant to one of the three main topics of the paper: green IT and carbon emissions, measuring data center energy efficiency, data center design and technology energy efficiency. And, each piece is considered a credible source based on authority, objectivity, quality of work, coverage of work and currency as outlined by Bell and Smith (2009).

Data Analysis Plan

Coding procedures. The eight coding steps of conceptual, or thematic, analysis described by Busch, et al. (2005) at the Colorado State University Writing Lab are used to perform the content analysis for this literature review. This is an interactive process and therefore key phrases and terms, relationships, and relevant content may evolve during data analysis. Steps are described below in relation to specific application to this study.

1. **Level of analysis**. Both words and phrases are coded.
2. **Pre-defined set of concepts and categories**. Only words and phrases relative to the main themes of the study are coded: data center energy measurement, data center energy consumption, data center design, data center best practices, green IT, data center emissions. Concepts and categories will evolve during the analysis phase of the study as they present themselves.
3. **Existence of a concept.** The fact that a concept exists in a work is of greater importance to this study than how frequently it is mentioned in the selected literature. A given concept is only coded once.
4. **Level of generalization**. Terms depicting similar concepts and categories are recorded as the same. For instance, “energy consumption” and “energy usage”, or “footprint” and “emissions” are treated as one. However, similar terms with different context such as “energy consumption” and “energy capacity”, or “modeling” and “monitoring” are coded separately.
5. **Translation rules.** Coding rules are used to ensure consistent categorization of terms. Each term is categorized once, and only once. Translation rules applied in this study are:

Sustainability = emissions	Data Center = server room	Financials=costs
Carbon Footprint	Facilities	Data Center budgets
Green House effect	Infrastructure	Corporate IT budgets
LEED	UPS	Projected cost increase
Carbon policy	HVAC	Cap & Trade
Green IT	CRAC	Implementation
	Cooling	
Energy = power	Design	
Policy	IT hardware	
Consumption = efficiency	Servers	
Monitoring	Storage	
Thermal modeling	Network	
Measuring = metrics	Classification	
Controlling	Sizing	
Thermal workload scheduling	Availability	

Table 2 - Translation Rules

6. **Irrelevant information**. Content that is clearly not related to this study in any way is excluded from further analysis. Content that does not impact context of coded information is also discarded.
7. **Code the texts**. Each literature piece is manually coded. A Microsoft Excel spreadsheet is used to track terms and phrases in each document. Relevant words and phrases are highlighted in printed copies of the literature. These highlighted items are then used to update the Excel spreadsheet entry for each document.
8. **Analyze results.** Once all documents have been coded an Entity-Relationship diagram (Chen, 1976) is used to show relationships amongst concepts, themes, and terms. Un-coded text might be reviewed for further insight.

Writing Plan

The Review of the Literature section of the paper is structured by a thematic rhetorical strategy, organizing research around ideas or insights (Literature Reviews, n.d.). The review is composed of three pre-determined themes, each relating to one of the three key subtopics of the study. Information is presented in a manner that reveals the development of each topic over time. The goal is to show how the field of data center management and sustainability is developing and to report the emerging best practices for reducing data center carbon emissions through energy usage efficiency gains. A preliminary thematic outline follows.

I. Green IT, carbon emissions, energy consumption

   a. Exploration of currently acceptable definitions of green IT, carbon emission, carbon footprint.

   b. Synthesis of why the green IT movement and energy consumption are relevant to established enterprise class data centers.

II. Measuring and managing data center energy efficiency

   a. Why data center efficiency and energy efficiency in a data center needs measurement and monitoring.

   b. Identification of emerging metrics and methodologies for energy efficiency measurements and reporting.

III. Data center design and technology energy efficiency

   a. Emerging design best practices.

   b. Emerging IT technological advances in energy efficiency.

Annotated Bibliography

This section provides a brief summary of key sources selected for the study. The abstracts included are either complete as published, or are slightly modified for length and/or content relevance. These abstracts provide a high level view of the study topic, lowering the data center footprint by increasing the efficiency of overall power consumption within the data center. The *Comments* sections describe how each reference is used in support of this study, and how credibility is established.

Avelar, V., Torell, W., & Hampel, T. (2008). *Data center baseline study report* [White paper] version]. Beaverton, OR.: The Green Grid.

Abstract. The two questions we sought to explore most deeply with this study were "what, if anything, is driving a need to increase energy efficiency", and "what are the obstacles to becoming more energy efficient in your data center." When the study was conceived, we did not expect results to contradict those of published surveys and articles. And it didn't. But neither did we anticipate it would challenge conventional wisdom. This, it did. The first question surfaced no unusual motivations for improved efficiency, with companies expressing concern with everything from impending regulation to increased energy costs to environmental issues. However, while specifics varied, every company we spoke with intended to use some, if not all, of the energy consumption that would be freed by greater efficiency to increase IT capacity. The question remains whether technology will be able to pick up the slack. The authors of this study believe that the IT and data center industries will make very significant improvements over the next few years. However, given accelerating technology adoption curves, we anticipate total energy consumption by data centers will increase.

Comments. The Green Grid is a highly active, and successful, global consortium of IT companies seeking to improve energy efficiency in data centers around the world. They are trying to create a standard set of metric, methods, and tools for measuring and managing energy usage within data centers. The organization includes the majority of large Information and Communications technology firms. The board of directors includes staff from: AMD, APC, Dell, EMC, HP, IBM, Intel, Microsoft, and Sun Microsystems. The Green Grid has successfully convinced the United States USEPA to use their popular metrics for Power Usage Effectiveness (PUE) in establishing PowerSTAR ratings for data centers.

The survey is relevant to this study because it provides insight into what steps, if any, a cross section of corporations plan to implement to address data center energy consumption issues. Information is used to support the discussion of data center design and technology energy efficiency, presented in the Review of Literature section of this paper. The study highlights possible emerging trends in the industry.

Battles B., Belleville C., Grabau S., & Maurier J. (2007). *Reducing datacenter power consumption through efficient storage* [White paper]. Sunnyvale, CA: Network Appliance, Inc,

Abstract. Exponential data growth is a reality for most data centers. Until recently, continuous improvements in price performance and $/GB have made it both easy and affordable to solve storage concerns simply by adding more disks to existing storage systems. However, IT executives are discovering that there are limits to that easy growth. Hitting any one of these limits significantly jeopardizes the ability of IT to meet the demands of business. Cooling is inextricably linked to power consumption. Every watt of power that enters the data center generates heat that must be removed from the environment—and to do so takes more power. As the heat increases, systems become more unstable and component failure rates rise. The cost of power to cool a system often is as much as powering the system itself. Environmental issues are gaining serious commercial momentum and, fueled by the growing number of local and global green initiatives, they are rising ever more insistently up the corporate agenda. More power-efficient storage solutions provide for business growth while saving power. Every watt of energy saved in the data center is a watt that is removed from an organization's carbon footprint and the global warming equation.

Comments. The white paper is relevant to this study because it discusses methods to increase energy efficiency of storage medium – a topic examined in the discussion of emerging IT technologies, presented in the Review of Literature section of this study.. The demand for storage in data centers is projected to continue expanding (Dogra, 2010) and therefore any improvement in energy usage of storage medium will have a significant return.

NetApp is a fortune 500 hundred company well recognized as an industry leader in storage hardware vendor arena. This specific white paper is incorporated in curriculum currently used by the Computer Science department of UC Berkley and cited by the USEPA's data center report to congress (2007).

Bodik, P., Armbrust, M., Canini, K. Fox, A., Jordan, M., & Patterson, D. (2008). *A case for adaptive datacenters to conserve energy and improve reliability* [Electronic version]. University of California at Berkeley.

Abstract. Although there is prior work on energy conservation in datacenters, we identify a new approach based on the synergy between virtual machines and statistical machine learning, and we observe that constrained energy conservation can improve hardware reliability. We give initial results on a cluster that reduces energy costs by a factor of 5, reduces integrated circuit failures by a factor of 1.6, and disk failures by a factor of 5. We propose research milestones to generalize our results and compare them with recent related work. Power is ranked as the #5 top concern of IT executives [Sca06], with availability and performance being #10 and #11 respectively.[1] This is not surprising given that powering and cooling a datacenter now rivals the cost of the hardware: each $1 spent on servers in 2005 required an additional $0.48 to power and cool it, expected to rise to $0.71 by 2010 [Sca06]. Also, new problems arise as power becomes the constraining resource in datacenters.

Comments. This technical research paper studies ways to improve energy consumption by combining the now common practice of server virtualization with statistical learning technologies. The implementation of statistical learning machines within commercial data centers is not imminent. However, the study is relevant to this research paper because of the significant efficiency gains achieved. Any methods of achieving such energy efficiency improvements within a data center may prove to be the birth of a new trend. The information is used to support the discussion of emerging best practices, presented in the Review of Literature section of this study.

The technical paper is published by the Computer Science department of UC Berkeley and cited in the USEPA’s report to Congress on data center energy efficiency (2007).

Ebbers, M., Galea, A., Schaefer, M., & Khiem, M. (2009). *The green data center: An idea for today.* Armonk. NY. : IBM

Abstract. The information technology (IT) industry is a significant user of resources as well as an enabler of efficiencies that reduce CO2 emissions. However, as companies continue to grow to meet the demands of their customers and as environmental concerns continue to be an issue, organizations are looking for ways to reduce corporate energy consumption and to become more environmentally responsible—in other words, to become *green*. What might surprise you is that often the most efficient, lowest cost solution is also the "greenest" one. So going green has caught on because it is a win/win effort. Today's business enterprises face many concerns in their data centers regarding how to achieve business sustainability. The cost of energy is growing, as is public awareness about data centers impacting our environment. Power, cooling, and space capacities are near the limit. With new environmental laws coming out and company images at stake, deploying environmentally friendly, or "green" solutions makes sense, from both cost and sustainability perspectives as well as a way to reduce organizational risk.

Comments. This book is a technical publication produced by IBM in the category commonly referred to as Red Books. They are developed by teams of professionals who volunteer their time to participate in residencies and workshops held by IBM on specific topics. These publications are the culmination of these residencies. The professionals are selected from applicants based on experience and references. The authors of this book come from Australia, France, Germany, and the United States. Two are IBM staff and two are customers. Each author has at least 20 years of experience in the industry.

The book is relevant to this literature review because it focuses on practical steps clients can take to "green" existing data centers. This book directly supports the discussion of green IT and carbon emissions, presented in the Review of Literature section of this paper.

Emerson. (2008). *Energy logic: Calculating and prioritizing your data center IT efficiency actions* [White paper]. Columbus, OH. : Emerson Electronics Co.

Abstract. Emerson Network Power's Energy Logic1 white paper provided a holistic, prioritized roadmap for reducing data center energy consumption. Energy Logic did not address data center efficiency directly because there is no universally accepted metric for data center output that could be used as the basis for this analysis. Now, the time has come for the industry to take the next step forward in achieving a true understanding of data center efficiency. The lack of true data center efficiency metric is challenging to IT and data center managers as they try to justify much needed IT investments to management. It also adds to the difficulty data center managers have in comparing efficiencies across their data centers to prioritize where efficiency-improving actions will have the greatest impact. In addition, they need to be able to track data center efficiencies over time. This paper not only shows how IT and data center managers can use an efficiency metric to address these challenges, but also provides a prioritized set of actions to gain the greatest improvement in efficiency.

Comments. Emerson Electronics Company is a multinational Fortune 500 company providing engineering services in industrial, commercial, and consumer markets. The authors of the paper are not listed so the credibility of the paper is based solely on the publisher. The paper is relevant to this study because it proposes an alternative set of

metrics to that put forth by The Green Grid. It supports the discussion of measuring and managing data center efficiency, presented in the Review of Literature section of this study. Because Emerson is also a significant presence in the technology market, this set of metrics and methodologies must be considered as possible trend setters.

Haas, J., Monroe, M., Pflueger, J., Pouchet, J., Snelling, P., Rawson, A. et al. (2009). *Proxy proposals for measuring data center productivity* [White paper]. Beaverton, OR: The Green Grid.

Abstract. This paper is a natural progression of the work that The Green Grid and the industry have already done to advance energy efficiency in data centers. A prior paper defined the Power Usage Effectiveness (PUE) and Data Center Infrastructure Efficiency (DCiE) metrics as follows: PUE = Total Facility Power / IT Equipment Power and its inverse, DCiE = IT Equipment Power x 100% / Total Facility Power. These metrics deal exclusively with power. A subsequent paper investigated this further and introduced the concept of a family of DCP metrics (referred to as DCxP metrics) based on a definition that measured useful work produced in a data center divided by any quantization of a resource consumed to produce that work. The paper also defined the first metric in this family, DCeP, where the resource consumed is energy. Thus, DCeP is defined as follows: DCeP = useful work produced in a data center/ total energy consumed in the data center to produce that work. The paper goes on to provide a detailed metric for accurately measuring useful work in a data center.

Comments. This white paper is the cornerstone for the rapidly growing trend within the data center industry of standardizing on the PUE as the metric of power efficiency. The white paper serves as the tipping point in the data center industry for how energy efficiency is defined and measured and supports the section on this same topic, presented in the Review of Literature section of this study. The paper is deemed credible because it is published by The Green Grid and the proposed methods of measurement have been at least partially adopted by the USEPA.

Hamann, H., Kessel, T., Lyengar, M., Chung, J., Hirt, W., Schappert, M., et al. (2009). Uncovering energy-efficiency opportunities in data centers . *IBM Journal of Research and development, 53*(3),10.1-10.12.

Abstract. The combination of rapidly increasing energy use of data centers (DCs), which is triggered by dramatic increases in IT (information technology) demands, and increases in energy costs and limited energy supplies has made the energy efficiency of DCs a central concern from both a cost and a sustainability perspective. This paper describes three important technology components that address the energy consumption in DCs. First, we present a mobile measurement technology (MMT) for optimizing the space and energy efficiency of DCs. The technology encompasses the interworking of an advanced metrology technique for rapid data collection at high spatial resolution and measurement-driven modeling techniques, enabling optimal adjustments of a DC environment within a target thermal envelope. Second, the static MMT measurements obtained at high spatial resolution are complemented by and integrated with a real-time sensor network. Third, an energy and thermal model analysis for a DC is presented that exploits both the high-spatial-resolution (but static) MMT data and the high-time-resolved (but sparse) sensor data. The combination of these two data types (static and dynamic), in conjunction with innovative modeling techniques, provides the basis for extending the MMT concept toward an interactive energy management solution.

Comments. This article is presented in the IBM research journal. All articles are peer reviewed. The authors are principle researchers at the IBM Austin Research Lab, are highly published and cited. The article is relevant to this literature review because it demonstrates new technological solutions to the data center energy efficiency problem under investigation by an influential research facility. This is a topic examined in the Review of Literature section of this study.

Hewlett-Packard. (2008). *Energy efficiency in the data center* [White paper]. Hewlett-Packard Development Company.

Abstract. Today's IT executives are not only expected to create and maintain high-availability IT environments, but they are also expected to implement green initiatives to satisfy customers, analysts, and government agencies that are worried about the impact of modern, energy-thirsty data centers on the environment.

Is such a dual mandate reasonable? Can companies be expected to maintain service levels and reduce their carbon footprints at the same time? This question is quickly becoming a moot point as events unfold before the industry. In July 2008 the UK government wrote to 10,000 businesses to warn them that they could be affected by the Carbon Reduction Commitment—a scheme that requires large businesses purchase yearly "carbon allowances" for their anticipated energy usage. Unused allowances can be redeemed at year's end, which is an incentive to reduce energy consumption.

In the following pages we will explore the drivers behind the exploding power demands of modern data centers, the current crop of solutions available to manage those demands, and the tremendous benefits—both economical and environmental—of taking steps to reduce those demands permanently.

Comments. This white paper is deemed credible because it is published by a key global server vendor providing another view on the problem of data center power usage. The paper is part of a body of such publications that help to establish possible industry trends. The paper supports the discussion of green IT and carbon emissions, presented in the Review of Literature section of this study.

Koomey, J. (2007). *Estimating total power consumption by servers in the U.S. and the World*. Oakland, CA: Lawrence Berkeley National laboratory, Analytics press.

Abstract. The amount of electricity used by servers and other Internet infrastructure has become an important issue in recent years as demands for new Internet services (like music downloads, video-on-demand, and Internet telephony) have become more widespread. One of the weaknesses in the literature on data center electricity use has been the lack of credible estimates of the aggregate power used by all servers and associated equipment in the U.S. and the world. The data on the floor area and power densities of data centers are anecdotal and limited by the proprietary nature of such data in most companies. Data on the installed base of servers are also closely held by the companies who track it, and server technology continues to change rapidly, necessitating constant updates to measurements of power used by particular server models.

This study estimates total electricity used by servers in the U.S. and the world by combining measured data and estimates of power used by the most popular servers with data on the server installed base. This study only assesses the direct electricity used by servers and associated infrastructure equipment. It does not attempt to estimate the effect of structural changes in the economy enabled by increased use of information technology, which in many cases can be substantial.

Comments. This is the landmark research publication, commissioned by the USEPA, to determine overall data center energy consumption nationally and globally. It is used as the cornerstone of the USEPAs data center report to congress published in 2007, and to drive Energy Star standards. This study is cited many times in professional, academic, and government publications. It has significantly influenced the nation's interest in data center energy consumption. This paper supports the discussion of the green IT and carbon emissions, presented in the Review of Literature section of this study.

Long, B., Freeman, J., & Patterson, M. (2009). *Assessment of EPA mid-tier data center at Potomac Yard* [White paper]. Beaverton, OR : The Green Grid.

Abstract. In April of 2008, The Green Grid signed a Memorandum of Understanding (MOU) with the Environmental Protection Agency (EPA) to promote energy efficiency in small to mid-sized data centers. This effort promotes the innovative efforts of the Information Technology (IT) industry and EPA to facilitate the improvements in the energy efficiency of computing facilities. According to the Environmental Protection Agency, data centers across the United States accounted for 1.5 percent of total US electricity demand in 2006 – equivalent to the annual electric consumption of the state of Florida – and have become one of the fastest growing users of energy. While the power consumed at these individual data centers may be small, they are numerous at EPA and other large organizations. IDC estimates that in the US there are 3.6 million servers in closets and rooms, 3.2 million mid-tier servers, and roughly 3.1 million servers in the enterprise-class space. Therefore, this report is relevant to roughly one-third of the servers in the United States. The following is the initial report documenting findings and making recommendations around a typical EPA mid-tier data center in the Washington

D.C. area: One Potomac Yard (OPY). Recommendations made in the report cover a wide range of specific and programmatic opportunities for midtier data centers across the US.

Comments. This study uses an USEPA data center as an example of the typical midsized data center. The Green Grid evaluates the data center for energy efficiency and makes recommendations for improvement. Any findings of this study are influential in the industry because it is conducted by The Green Grid and therefore relevant to this literature review. This publication supports the discussion of emerging trends and best practices, presented in the Review of Literature section of this study.

Mathew, P., Greenberg, P., Ganguly, S., Sartor, D., & Tschudi, P. (2009). *How does your data center measure up? Energy efficiency metrics and benchmarks for data center infrastructure systems.* Lawrence Berkeley National Laboratory.

Abstract. Data centers are among the most energy intensive types of facilities, and they are growing dramatically in terms of size and intensity [USEPA 2007]. As a result, in the last few years there has been increasing interest from stakeholders - ranging from data center managers to policy makers - to improve the energy efficiency of data centers, and there are several industry and government organizations that have developed tools, guidelines, and training programs. There are many opportunities to reduce energy use in data centers and benchmarking studies reveal a wide range of efficiency practices. Benchmarking is an effective way to compare one facility to another, and also to track the performance of a given facility over time. This article presents the key metrics that facility managers can use to assess, track, and manage the efficiency of the infrastructure systems in data centers, and thereby identify potential efficiency actions. Most of the benchmarking data presented in this article are drawn from the data center benchmarking database at Lawrence Berkeley National Laboratory (LBNL). The database was developed from studies commissioned by the California Energy Commission, Pacific Gas and Electric Co., the U.S. Department of Energy and the New York State Energy Research and Development Authority.

Comments. This article shows how benchmarking can be used to help data center managers compare and track energy efficiency performance over time. The authors are academic researchers in the same lab that was initially commissioned by the USEPA to study data center power consumption and has a significant database of data center usage statistics. Recommendations in the paper must be considered as significant influencers of industry best practices and trends, because it was produced by an academic laboratory well established as an expert in data center energy efficiency research,

Moore, J., Chase, J., Farkas, P. & Ranganathan, P. (2005). *Data center workload monitoring, analysis, and emulation*. Duke University.

Abstract. Over the last ten years we have witnessed a shift from large mainframe computing to commodity, off-the-shelf clusters of servers. Today's data centers contain thousands or tens of thousands of servers, providing services and computation for tens or hundreds of thousands of users. In addition to traditional IT challenges such as server management, security, and performance, data center owners now must deal with power and thermal issues, previously the domain of facilities management. These trends will continue to accelerate as organizations acquire bladed servers and consolidate multiple, smaller clusters into centrally-located data centers. However, in spite of these trends, there has been no corresponding change in emphasis in the methods and toolkits that

target system instrumentation, analysis, management, replay, and emulation. This paper seeks to address this gap. We focus on methods and toolkits to enable the automated collection and analysis of workload traces from data centers, and use those traces as the basis for repeatable and verifiable experiments and workload emulation.

Comments. This paper is both heavily referenced, and includes many academic references. It is published by Duke University and is co-authored by two members of Duke University's computer science department and two researchers at the Hewlett Packard Labs. The concept of data center workload monitoring is significant to this literature review because data center efficiency is a function of workload, and data center power consumption is also a function of workload. Data center power usage efficiency cannot be measured if workload is unknown. The paper supports the discussion of measuring and managing data center efficiency, presented in the Review of Literature section of this study.

Rasmussen, N. (2006). *Electrical efficiency modeling of data centers* [White paper]. West Kingston, RI : American Power Conversion.

Abstract. Conventional models for estimating electrical efficiency of data centers are grossly inaccurate for real-world installations. Estimates of electrical losses are typically made by summing the inefficiencies of various electrical devices, such as power and cooling equipment. This paper shows that the values commonly used for estimating equipment inefficiency are quite inaccurate. A simple, more accurate efficiency model is described that provides a rational basis to identify and quantify waste in power and cooling equipment. The 10-year total cost of ownership (TCO) for network-critical

physical infrastructure (NCPI) in a typical data center can be $80,000 to $150,000 per rack. Of this cost, the cost of electrical power consumption is a significant contributor, on the order of 20% of the total cost1. This is of interest because much of the electrical power consumption is wasted (in the form of heat energy) and a significant amount of this waste is avoidable. It is estimated that, worldwide, data centers consume 40,000,000,000 kW-hr of electricity per year and the reduction of waste associated with this consumption is a significant public policy issue as well as a major financial concern to data center operators.

Comments. This article is referenced by a number of white papers published by The Green Grid, as well as the USEPA data center report to congress published in 2007. It is significant to this study because it provides an overview of how a data center's electrical usage can be modeled, and then how the model can be applied to improve efficiency. Most importantly it provides what is becoming a commonly accepted definition of what data center efficiency is. The paper supports the discussion of measuring and managing data center efficiency, presented in the Review of Literature section of this study.

Radicke, J., Roden, B., & Yunke, L. (2009). *Exploring views on data center power consumption and server utilization*. Lund University.

Abstract. The primary purpose of this Thesis is to explore views on Green IT/Computing and how it relates to Server Virtualization, in particular for Data Centre IT environments. Our secondary purpose is to explore other important aspects of Server Virtualization, in the same context. The primary research question was to determine if Data Centre (DC) power consumption reduction is related to, or perceived as, a success factor for implementing and deploying server virtualization for consolidation purposes, and if not, what other decision areas affect Server Virtualization and power consumption reduction, respectively. The conclusions from our research are that there is a difference of opinion regarding how to factor power consumption reduction from server equipment, both from promoters and deployers. However, it was a common view that power consumption reduction was usually achieved, but not necessarily considered, and thus not evaluated, as a success factor, nor that actual power consumption was measured or monitored after server virtualization deployment. We found that other factors seemed more important, such as lower cost through higher physical machine utilization, simplified high availability and disaster recovery capabilities.

Comments. This master thesis paper is included in the main body of literature to be reviewed because it is the only academic literature review identified on the topic of data center energy consumption. It is significant in that it shows prior work that compliments this study. This paper supports the discussions of Green IT and carbon emissions, and the emerging trends, presented in the Review of Literature section and the emerging trends section of this study.

Rawson, A., Pfleuger, J., & Cader, T. (2008). *Green Grid data center power efficiency metrics: PUE and DCIE* [White paper]. Beaverton, OR : The Green Grid.

Abstract. The Green Grid is an association of IT professionals seeking to dramatically raise the energy efficiency of datacenters through a series of short-term and long-term proposals. This is an update to the very first white paper published by The Green Grid in February 2007 called “Green Grid Metrics: Describing Data Center Power Efficiency” to refine the nomenclature and intent of that paper. In that paper, The Green Grid proposed

the use of Power Usage Effectiveness (PUE) and its reciprocal, Datacenter Efficiency (DCE) metrics, which enable datacenter operators to quickly estimate the energy efficiency of their datacenters, compare the results against other datacenters, and determine if any energy efficiency improvements need to be made. Since then PUE has received broad adoption in the industry but DCE has had limited success due to the misconception of what data center efficiency really means. To promote these metrics and drive greater datacenter energy efficiency for businesses around the world, The Green Grid will publish future white papers that provide detailed guidance on using these metrics. We will also continue to collaborate with organizations such as the USEPA, ECMA and Climate Savers that promote a similar goal and vision.

Comments. This white paper is a key publication outlining the metrics recently adopted by the USEPA for the measurement of power usage effectiveness and efficiency within a data center. It is highly probable that these metrics are going to become the industry standard for measuring data center energy consumption and efficiency. The paper supports the discussion of measuring and managing data center efficiency, presented in the Review of Literature section of this study.

Seaton, I. (2004). *Best practices & concepts for computer room cooling* [White paper]. Westlake Village, CA: Chatworth Products,Inc.

Abstract. Computer equipment cooling air delivered from a raised floor air source can be optimized by applying some basic thermal management concepts, which sometimes contradict conventional wisdom or intuition. This optimization can either produce increased cooling or reduced cooling costs, and often both apparently mutually exclusive objectives can be met simultaneously. This paper will review the effects of heat on

computer equipment and some of the benefits of improving cooling based on some simple principles of conduction and forced convection heat transfer and an understanding of fluid dynamics both under the raised floor and within an equipment enclosure. While historical practice and conventional wisdom have typically focused on getting the heat out of the cabinet and fretting over hot spots (usually somewhere around the upper rear of the cabinet), the principles of forced air convection heat transfer indicate clearly that the real benefit to CPU performance and life results from the delivery of the coldest available air directly to the conductive surface area attached to the microprocessor. In fact, in the dynamics of heat transfer, the temperature of the delivered air (Tf) is the only variable over which the equipment installers and users have any control and is the objective of air movement management.

Comments. This publication is included in the coding set for data analysis, because it provides a comprehensive, high level, outline of best practices in 2004. It serves as a base line for current trends. The publisher is a leading manufacturer of data center infrastructure components that has already adopted The Green Grid's metrics as standard measurement of energy efficient. This specific white paper is sited in a number of other vendor publications. The paper supports the discussion of data center design and technology trends, presented in the Review of Literature section of this study.

Stanley, J. Brill, K., & Koomey, J. (2006). *Four metrics define data center greenness* [White paper]. Sanata Fe, NM. Uptime Institute.

Abstract. 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 of 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 user to identify profitable efficiency opportunities. Acting on those opportunities can potentially save large amounts of money across a data center portfolio by deferring or possibly elimination major new capital investment.

Comments. The Uptime Institute is a well established and respected organization specializing in improving datacenter availability and reliability. It is a provider of educational and consulting services to Facilities and Information Technology organizations. The paper is co-authored by a Ph.D. candidate in the Energy and Resources Group at UC Berkeley, the founder of the Uptime Institute, and a scientist at the Lawrence Berkeley Nation Lab who is also a consulting professor at Stanford University.

This white paper is significant because it provides metrics for overall "greenness" of a datacenter, not just for the measurement of power efficiency usage. It is also significant because the authors are prominent figures in the industry of data center management and energy usage, and thus considered influential. This paper supports the discussion of Green IT and carbon emissions, as well as the measuring and managing data center efficiency, presented in the Review of Literature section of this study.

Symantec. (2009). *Green IT Global results: Survey results May 2009*. Mountain View, CA.: Symantec Corporation.

Abstract. Green used to be associated primarily with the environment. Today it also refers to saving money. In this year's Green IT Report, companies surveyed reported that for them, Green IT means lowering electricity costs as well as sustainable energy practices. In this year's annual Green IT Survey, Symantec has found that Green IT has reached critical mass. Virtually all the companies we surveyed (97 percent) are discussing their green strategy, with just two percent saying it is unimportant.

They are not just talking, either. Green budgets are on the rise and IT is more than willing to pay a premium on energy efficient products—as much as 20 percent more according to 41 percent of the respondents. Finally, IT finds itself leading the way in corporate green policy, which is not surprising given the important role that it plays in a company's operation and the price organizations find themselves paying for energy. In addition to having a seat at the corporate green table, IT is also involved in a bevy of Green IT initiatives.

Comments. This set of survey results is published by a well established information technology vendor. Each regional survey publication includes the methodology, the demographic of respondents, the individual questions, and a description of resulting data sets. The survey is significant in its currency and topic. The majority of questions address energy efficiency and/or data center efficiency. The survey supports the discussions of Green IT and carbon emissions, as well as the emerging trends, presented in the Review of Literature section of this study.

Torres, J., Carrera, D., Hogan, K., Gavalda, R., Betran, V., & Poggi, N. (2008). *Reducing wasted resources to help achieve green data centers.* Technical University of Catalonia.

Abstract. In this paper we introduce a new approach to the consolidation strategy of a data center that allows an important reduction in the amount of active nodes required to process a heterogeneous workload without degrading the offered service level. This article reflects and demonstrates that consolidation of dynamic workloads does not end with virtualization. If energy-efficiency is pursued, the workloads can be consolidated even more using two techniques, memory compression and request discrimination, which were separately studied and validated in previous work and are now to be combined in a joint effort. We evaluate the approach using a representative workload scenario composed of numerical applications and a real workload obtained from a top national travel website. Our results indicate that an important improvement can be achieved using 20% less servers to do the same work. We believe that this serves as an illustrative example of a new way of management: tailoring the resources to meet high level energy efficiency goals.

Comments. This academic research publication is included in the IEEE database of articles and includes a very impress list of references and thus considered to be a credible piece of literature. It is included in this review as an example of emerging academic research on the problem of greening data centers. It proposes a novel method of server and workload consolidation within existing data centers. The paper supports the discussion of data center design and technology efficiency, presented in the Review of Literature section of this study.

United States Environmental Protection Agency. (2007). *Report to congress on server and data center energy efficiency.* Prepared for the U.S. Environmental Protection Agency, Energy Star program, by Lawrence Berkeley National Laboratory.

Abstract. The United States (U.S.) Environmental Protection Agency (USEPA) developed this report in response to the request from Congress stated in Public Law 109-431. This report assesses current trends in energy use and energy costs of data centers and servers in the U.S. and outlines existing and emerging opportunities for improved energy efficiency. It provides particular information on the costs of data centers and servers to the federal government and opportunities for reducing those costs through improved efficiency. It also makes recommendations for pursuing these energy-efficiency opportunities broadly across the country through the use of information and incentive-based programs. As our economy shifts from paper-based to digital information management, data centers — facilities that primarily contain electronic equipment used for data processing, data storage, and communications networking — have become common and essential to the functioning of business, communications, academic, and governmental systems. Data centers are found in nearly every sector of the economy: financial services, media, high-tech, universities, government institutions, and many others use and operate data centers to aid business processes, information management, and communications functions.

Comments. This is the milestone USEPA report to congress that raised the visibility of data center energy usage nationally. It is very significant to this study because it is highly influences the surging global interest in the problem of energy consumption and efficiencies within data centers. The report directly supports the discussion of Green IT and carbon emissions, presented in the Review of Literature section of this study.

United States Environmental Protection Agency. (August 2009). *Methodology for incorporating source energy use*.

Abstract. USEPA's national energy performance ratings evaluate the performance of buildings that use all types of energy. To compare this diverse set of commercial buildings equitably, the ratings must express the consumption of each type of energy in a single common unit. USEPA has determined that source energy is the most equitable unit of evaluation. Source energy represents the total amount of raw fuel that is required to operate the building. It incorporates all transmission, delivery, and production losses, thereby enabling a complete assessment of energy efficiency in a building. Site energy may be delivered to a facility in one of two forms: primary and/or secondary energy. Primary energy is the raw fuel that is burned to create heat and electricity, such as natural gas or fuel oil used in onsite generation. Secondary energy is the energy product (heat or electricity) created from a raw fuel, such as electricity purchased from the grid or heat received from a district steam system. A unit of primary and a unit of secondary energy consumed at the site are not directly comparable because one represents a raw fuel while the other represents a converted fuel.

Comments. This document provides technical details on the methodology used by the USEPA to include source energy in the Energy Star performance ratings for energy consumption of buildings. It is significant since this same measurement will be applied to stand alone data center facilities. The report support the discussion of measuring and managing data center efficiency, presented in the Review of Literature section of this study.

Wiedmann, T., & Minx, J.(2007). *A definition of 'Carbon Footprint'* [Electronic version]. ISA uk Research & Consulting.

Abstract. The term ‘carbon footprint’ has become tremendously popular over the last few years and is now in widespread use across the media – at least in the United Kingdom. With climate change high up on the political and corporate agenda, carbon footprint calculations are in strong demand. Numerous approaches have been proposed to provide estimates, ranging from basic online calculators to sophisticated life-cycle analysis or input-output-based methods and tools. Despite its ubiquitous use however, there is an apparent lack of academic definitions of what exactly a ‘carbon footprint’ is meant to be. The scientific literature is surprisingly void of clarifications, despite the fact that countless studies in energy and ecological economics that could have claimed to measure a ‘carbon footprint’ have been published over decades. This report explores the apparent discrepancy between public and academic use of the term ‘carbon footprint’ and suggests a scientific definition based on commonly accepted accounting principles and modeling approaches. It addresses methodological questions such as system boundaries, completeness, comprehensiveness, units, and the robustness of the indicator.

Comments. This research paper is the addresses the lack of a commonly accepted definition of a ‘carbon footprint’ and provides one of its own, as noted in the Purpose section of this paper. The paper is heavily referenced by a broad range of publications both academic and professional. Of the two authors, one is a member of the publishing consulting group and the other is on staff at the University of York. Their paper includes a significant number of citations. All indications are that this report can be considered to be a credible source of information. The paper also directly supports the discussion of Green IT and carbon emissions, presented in the Review of Literature section of this study.

Review of the Literature

This section of the study provides a review of selected publications on the topic of green IT, data centers, and energy efficiency within data centers. The primary purpose of the literature review is to highlight the latest research and emerging themes in the industry regarding: (a) green IT and carbon emissions, (b) measuring and managing data center power consumption, and (c) data center design and technology improvement as applicable to established enterprise data centers. The section is presented in three parts, one each for the themes listed above.

Theme I: Green IT and carbon emissions

Analysis of selected literature shows minimal discussion of data centers and carbon emissions prior to 2004. There is a sharp and continual increase in literature beginning 2005 with no indication of having reached a constant state yet. The proliferation of literature on the topics of green IT and carbon emissions, data center energy consumption, and data center design reflects the momentum that environmental issues are gaining in corporate agendas (Battles, Belleville, Grabau, & Maurier, 2007; Gonella, Lee, Chen, & Tian, 2009).

The concept of green computing, or green IT began in 1992 when the US Environmental Protection Agency (USEPA) initiated the Energy Star program. The purpose of Energy Star is to provide for volunteer energy-efficiency citification (Baroudi, Hill, Reinhold & Senxian, 2009). Currently much electronic equipment includes Energy Star compliance in their product description (Ruth, 2009). A number of factors have converged since 1992 to raise green IT and data center energy consumption to the forefront of many IT executives awareness (Bodik, Armbrust, Canini, Fox, Jordan & Patterson, 2008). Factors include new local and global environmental regulations (Ebbers, Galea, Schaefer, & Khiem, 2009; O'Flynn, 2008), power capacity constraints (Fan, Weber, & Barroso, 2007; Kaplan, Forrest, & Kindler, 2008; Nexsan, 2009; PG & E, 2006), increased energy costs(Gartner Inc, April 2007), and a greater focus on corporate social responsibility(Ebbers, Galea, Schaefer, & Khiem, 2009; Gartner Inc, April 2007; Ohara, 2008).

In its simplest definition, green IT is finding ways to process and store information while minimizing carbon footprint (Baroudi, Hill, Reinhold, & Senxian, 2009). A carbon footprint is a somewhat nebulous term that is commonly used to refer to the amount of gaseous emissions produced by human consumption and production that influence climate change, such as carbon dioxide(CO_2) (Wiedmann & Minx, 2008). The burning of fossil fuels, such as oil and coal for applications like power generation, has been directly linked to increased CO_2 emissions (United States Government, 2008). A 2007 USEPA report to congress claims that data centers accounted for 1.5% of the total U.S. electricity consumption, which has been translated by some to be 2% of national $C0_2$ emissions (Ebbers, Galea, Schaefer & Khiem, 2009; Koomey, 2007). To make matters worse, the USEPA also states in this same report that by 2011, U.S. data centers are projected to increase annual consumption by an amount roughly equivalent to the output of 10 new power generation plants (USEPA, 2007).

As global interest in climate change and environmental stewardship has increased, environmental laws addressing carbon emissions are beginning to affect the IT industry (Emerson, 2008). In particular, there is increasing pressure to reduce the IT carbon footprint (O'Flynn, 2008; Ohara, 2008). For example, the USEPA is continuing to analyze data center energy consumption, the European Union has established a directive to decrease energy consumption 20% by 2020, and Australia requires high energy consumer businesses to prepare assessment and action plans (Ebbers, Galea, Schaefer, & Khiem, 2009; Mainline, 2009).

Increases in computer speed and capacity have been achieved through increased processor power and chip density. The result has been a correlated growth in power consumption and heat generation within data centers (Nexsan, 2009; PG & E, 2006). One study published in 2006 found that data center use up to 40 times the amount of energy as conventional office building (Greenberg, Mills & Tschudi, 2006). The associated cost of powering and cooling of a data center is also rapidly increasing (AMD, 2005; Bodik, Armbrust, Canini, Fox, Jordan & Patterson, 2008). By 2011, power costs are projected to be an additional 70 cents, or 70%, in ongoing operational costs for each dollar spent on hardware (Fan, Weber, & Barroso, 2007).

A number of relatively recent IT industry surveys indicate that a significant percentage of organizations consider data center power consumption and cooling an issue of concern (AMD, n.d; Bodik, Armbrust, Canini, Fox, Jordan & Patterson, 2008; Gartner, April 2007). Power consumption is fast becoming a constraining resource in data centers, and many are finding that they are very close to being out of site capacity (Bodik, Armbrust, Canini, Fox, Jordan & Patterson, 2008; Brill, 2006). And the ever increasing demand on the national power grid by data centers has caused the USEPA to express concern about the ability of national transmission systems to service power in some areas (USEPA, 2007).

As public awareness of climate change and environmental issues expands, pressure is mounting on the IT industry to take proactive measures to improve efficiencies in order to minimize their corporate carbon footprint. For some companies, there is a growing awareness of their environmental profile and the potential impacts to public relations (Mainline, 2009; Ohara, 2008; Gartner, April 2007). In response, more corporations are incorporating environmental stewardship in fundamental pillars or goals (Gonella, Lee, Chen & Tian, 2009).

Theme II: Measuring and managing data center efficiency

Data center managers are under pressure not only to decrease total cost of ownership (AMD, 2005; Moore, Chase, Farkas & Ranganathan 2005; Microsoft, 2008), but also to reduce corporate energy consumption and to become more environmentally responsible -- that is to become green. (Ebbers, Galea, Schaefer, & Khiem, 2009; Emerson, 2008; Gartner, April 2007; Ohara, 2008). At the same time, data center energy consumption has drastically increased across every business sector (Hamann, Kessel, Lyengar, Chung, Hirt, Schappert, et al. 2009). In order for a data center to decrease its carbon footprint and contain costs, it must increase effective power usage (Ohara, 2008; SVLG Intel, 2009). The first step to increasing energy efficiency is to assess how the data center is using power (Fan, Weber, & Barroso, 2007). Surprisingly, in 2009 only 30% of organizations in the United States were actively investigating data center power consumption and cooling (AMD, n.d).

Metrics. Standards for green IT are in early development (Baroudi, Hill, Reinhold & Senzian, 2009). However, various industry organizations, government bodies, and academics have begun researching means of measuring data center power efficiency and consumption, or greenness (Hamann, Kessel, Lyengar, Chung, Hirt, Schappert, et al. 2009). The key questions being asked are how to define IT equipment energy efficiency, and how does this equipment compare (Ebbers, Galea, Schaefer, & Khiem, 2009).

A number of metrics have been proposed to assess and track data center greenness (Emerson, 2008; Haas, et al., 2009; Hewlett-Packard, 2008; Stanley, Brill & Koomey 2006). A 2006 Uptime Institute white paper is an early attempt at defining metrics. It separates data center energy consumption into two components: IT productivity per watt component, and site infrastructure energy efficiency (Brill, 2006). One global non-profit trade organization of IT professionals, The Greed Grid, has been exceptionally active in this area of green IT. Its mission statement is to provide recommendation on best practices, metrics, and technologies designed to improve overall data center efficiency (Rawson, Pfleuger, & Cader, 2008).

In 2007, The Green Grid introduced two metrics for data center efficiencies with the primary intention of providing established data centers a method to measure, track, and ultimately optimize efficiency. The two metrics are Power Usage Effectiveness (PUE) and Data center Efficiency (DCE). The PUE is total facility power usage divided by IT equipment power usage. The PUE quickly gained traction, however the DCE did not. In 2008 the organization refined its DCE metric proposing the Datacenter Infrastructure Efficiency (DCiE) as being the reciprocal of the PUI. That is, the DCiE = 1/PUE= IT equipment power/Total facility Power x 100% (Rawson, Pfleuger, & Cader, 2008). The Green Grid provides the following useful definitions:

- IT equipment power is defined as the energy load associated with any IT equipment within the data center including primary computing equipment – processors, storage, network – and secondary equipment such as switches, monitors, and workstations used to control the data center (Rawson, Pfleuger, & Cader, 2008).
- Total facility power is defined as the energy load associated with IT equipment and everything that supports the IT equipment. This includes power delivery components, cooling systems, data center lighting, etc. (Rawson, Pfleuger, & Cader, 2008).

The DCiE is reciprocal of the PUE expressed as percentage. For example, a PUE of 2.0 has a DCiE of 50% indicating that half the power consumed by the data center is going to the servers and other IT equipment. This number does not indicate how well that power is consumed, that is the efficiency of the equipment, but only where it is consumed (Baroudi, Hill, Reinhold &

Senxian, 2009). IBM published a study in 2009 indicating the average DCiE is 44%, meaning that only 44% of the total energy fed into the average data center powers IT equipment. IBM also states an excellent DCiE is more than 60%, showing clear room for industry improvement overall (Ebbers, Galea, Schaefer, & Khiem. 2009).

In 2008 another proxy measurement for computing output, Compute Units per Second (CUPS), was introduced by the electronics industry manufacturing giant Emerson to measure server efficiency. One MegaCUPS is set to the average server performance in 2002. A server draws power in hundreds of watts, thus MegaCUPS allows measurements in CUPS/watt to be expressed as an integer. The metric is used to identify and track efficiencies of each data center location, rather than a holistic data center view. Emerson (2008) claims applying this measurement shows that while data center energy consumption has increased significantly, the gains in output outweigh the cost of these increases. NetApp, a storage solution vendor, has introduced the concept of measuring storage server efficiency as watts consumed per usable Terrabyte (Battles, Belleville, Grabau & Maurier, 2007). With the continued growth of storage requirements in the data center, storage efficiencies are gaining special attention as well (Dogra, 2010).

In early 2009, the USEPA announced it would devise its own metric, Energy Usage Effectiveness (EUE), rather than apply the now commonly used PUE to benchmark data center energy efficiency (United States Environmental Protection Agency, September 2009). However, the agency later reverted to using the PUE due to strong industry feedback (United States Environmental Protection Agency, November 2009).

Interestingly, in 2008 the U.S. Department of Energy (DOE) published a pamphlet titled *A Quick Start to Energy Efficiency* detailing the usage of both the PUE and the DCiE as high level facility metrics. This brochure states that improving PUE from 2.0 (50% DCiE) to 1.6 (62.5%) for a data center with a 2.5 MW load will result in 20% energy savings and over $800,000 in annual savings.

The Green Grid believes that a true data center efficiency rating must include the productivity of the data center (Haas, Monroe, Pflueger, Pouhet, Snelling, Rawson et al, 2009). They are currently working on another metric, DCeP, for the efficiency of IT equipment. It will measure productivity by work done per unit of energy (Mathew, Greenberg, Ganguly, Sartor, & Tschudi, 2009). The purpose of the DCeP is to quantify useful work over a period of time compared to the energy it requires. The metric should be calculated for an individual IT device or a cluster of computing equipment.

Methodology. As with data center efficiency metrics, there is not yet a clear industry framework, or roadmap, for applying efficiency metrics or reducing data center energy consumption (Emerson, 2008). It is recognized in all of the selected literature addressing this theme that in order to reduce energy consumption in a data center, how the energy is used has to be studied and understood. Power usage cannot be managed or controlled if it is not first measured (Baroudi, Hill, Reinhold & Senxian, 2009; Bouley & Brey, 2009). Benchmarking the energy consumption of a data center is the recommended first step. It provides a baseline measurement that can be used to determine change in power usage over time (Baroudi, Hill, Reinhold & Senxian, 2009; Greenberg, Mills & Tschudi, 2006). The implementation of metric driven management of data center energy efficiency by itself has shown increased efficiency of 10% (Hamann, Kessel, Lyengar, Chung, Hirt, Schappert et al, 2009). For further efficiency gains, the physical infrastructure of the data center has to be analyzed to identify potential energy losses (Bouley & Brey, 2009).

Much of the published literature reviewed for this study uses The Green Grid set of metrics to measure and track overall data center efficiencies and power usage (Baroudi, Hill, Reinhold & Senxian, 2009; Ebbers, Galea, Schaefer & Khiem, 2009; Mathew, Greenberg, Ganguly, Sartor & Tschudi, 2009). However, there are a number of frameworks, tools, and approaches currently marketed by vendors for the ongoing monitoring and management of data center power consumption. IBM has announced a combination of software and framework called the *Green Transformation Workbench,* which is intended to help assess and identify enhancement opportunities for data center efficiency. Gonella, Lee, Chen and Tian (2009) describe this framework as one that uses a holistic approach that includes process, people and infrastructure. Emerson (2008) focuses on measuring and managing at the server level through a framework it is marketing as *Energy Logic*. The McKinsey group advocates for their CAFÉ and CADE measurements and framework (Kaplan, Forrest & Kindler, 2008).

Theme III: Data center design and technologies

Typically power supplied to the data center is split, using a power-switching mechanism, into two paths: one for IT equipment and one for systems that support the IT equipment. The power path for IT equipment is maintained by uninterruptible power supplies (UPS) and distributed by power distribution units (PDU) that then supply the IT equipment (Hamann, Kessel, Lyengar, Chung, Hirt, Schappert, M., et al. 2009). Eventually, all power going into the IT equipment is converted to the unintended by-product, heat. To prevent degradation of the processor chips within the equipment, the heat must be removed from the chips and the systems.

The greatest opportunity for environmental improvement for many IT organizations is in data center cooling (Baroudi, Hill, Reinhold & Senxian. 2009). It is estimated that every watt saved in powering the IT equipment results in another watt saved in heat load. And, these savings also impact the UPS and cooling systems potentially reducing overall consumption twofold (Ebbers, Galea, Schaefer, & Khiem, 2009).

All power used by IT equipment is converted to heat that then must be removed from the environment, which in turn uses more power (Freeman, 2009). Most of the equipment is air-cooled, and a typical enterprise data center contains thousands of IT devices resulting in complex hot airflow paths. All this excess heat is removed by air conditioning systems (Dunlap & Rasmussen, 2006). Historically data centers are cooled on a room based method; that is, one or more computer room air conditioning (CRAC) systems pushed cold air into the room while removing warmer air. The air conditioners also mixed the air within the room to create an average temperature throughout preventing any hot spots. Unfortunately, this method of cooling is only effective when the power to run the air conditioning equipment is a small percentage of the power consumed by the data center (Greenberg, Mills & Tschudi. 2006). As the number of processors on single chips continues to grow, the energy consumption by these chips rises along with the heat produced (Baroudi, Hill, Reinhold & Senxian. 2009). And, as the computing power that can be fit into a given space increases, room-based cooling is less effective. (Dunlap & Rasmussen, 2006)

Row-oriented architecture. Hot and cold isle layout has become a commonly implemented design for server racks and other computing equipment in a data center. The goal of a hot aisle/cold aisle configuration is to conserve energy and lower cooling costs by managing air flow. This design involves lining up server racks in alternating rows with cold air intakes at the front of the racks facing one way, and hot air exhausts from the rear of the rack facing the other. Some designs isolate hot and cold isles using various types of containment systems (AMD, 2006; Bouley & Brey, 2009; Microsoft, 2008; Seaton, 2004). Raised floors, already utilized in most data centers, are used to send cold air to the cold isles. Air ducts are then used to send warm air from the ceiling above the hot isles back down to the CRACs. For added efficiency CRAC units can be dedicated to each row. By mounting units under the floor and overhead, airflow paths can be significantly shortened. Short paths in turn mean that the fans within the CRACs have a reduced workload, requiring less power, therefore increasing energy efficiency. This reduction is significant, because in some data centers the fan power consumption can be greater than the total IT equipment energy usage (Dunlap & Rasmussen, 2006).

Rack-oriented architecture. In this design, CRAC units are dedicated to a single rack of IT equipment. The units are typically either mounted directly to a rack or housed within it (Ebbers, Galea, Schaefer & Khiem, 2009; Dunlap & Rasmussen, 2006). By placing the units directly in the racks the airflow paths are shorter, exactly defined, and immune to room constraints or variations. This architecture also enables the use of CRAC units designed for targeted needs such as varying power densities amongst the racks. (Dunlap & Rasmussen, 2006).

Virtualization. Another much discussed design trend is server virtualization (Baroudi, Hill, Reinhold & Senxian, 2009; Loeffler, 2009). This design places multiple, self contained, logical servers on a single physical server. It is possible for many virtual servers to share one physical server because the average physical server is severely underutilized. Physical server utilization rates ranging from 10% to 25% are commonly quoted (ASE, 2009; Loeffler, 2009; Loper & Parr, 2007). Virtualization can increase energy efficiency by minimizing the amount of IT equipment required in the data center. An idle server will consume 50% of the power used by a fully utilized server. Therefore, while additional workload on a given server will increase the power consumed by that server, there will be a net decrease in over power consumption and heat generation within the data center though server virtualization (Loeffler, 2009).

IT equipment. IT equipment vendors are taking cooling one step farther by delivering water cooling devices that serve as the rear door of servers and with direct water cooling at the chip level (Ebbers, Galea, Schaefer & Khiem, 2009). Storage vendors are decreasing storage space requirements through new data deduplication technology. Various deduplication technologies can decrease the total amount of storage required in the average enterprise data center by 50% (Freeman, 2009). Software vendors are offering power management tools. These software products manage electrical systems of all sizes from sets of buildings to individual PCs. The Energy Star Portfolio Manager product reports potential energy saving of 6% to 15% in over 20 buildings located in the Pacific Northwest (Ruth, 2009). Yale University reported energy savings of $40 per PC and Verizon predicted $7 million savings for 185,000 PCs (Ruth, 2009).

Conclusion

There is an increased emphasis on data center efficiency in literature published at the end of this decade, reflecting expanded research into the problem by both industry and academic sectors. A number of trends and best practices emerge from the exploration of the selected works.

Emerging Trends

Green computing. According to recent surveys, companies are increasingly assessing the impact of their IT needs on the environment. Many have incorporated green criteria in their procurement processes (Mainline, 2009; Radicke, Roden & Yunke, 2009). However, there are still a number of hurdles facing the IT industry, including: (a) difficulty in accurately measuring data center efficiencies, (b) a lack of clear standardized roadmaps from equipment vendors on efficiency gains in future produces and finally, and (c) a lack of industry best practices and tools for managing power consumption (Kumar, 2009; Avelar, Torell & Hampel, 2009).

Climate change and carbon footprint are getting more attention at the global, national, and local levels. As government agencies are exploring ways to regulate energy use (Ohara, 2008) green principles are receiving more attention within IT organizations (Hewlett-Packard, 2009). Historically, the criticality of the loads supported by enterprise data centers has made reliability and power capacity the primary data center design criteria. Power efficiency and life cycle costs are not of the same priority (PG & E, 2006). According to The Green Grid, there are basic improvements that many data centers have yet to take such as hot and cold aisles arrangements and inventorying infrastructure to discover devices that are not used but still powered up (Avelar, Torell, & Hampel, 2009). Becoming green and reducing carbon emissions is a long-term project for data centers and the first steps highlighted in this study need not wait for industry cohesion on measurements and frameworks. As noted in this study, small incremental steps taken now will start the thoughtful marathon of continuous efficiency management, resulting in financial gains and ultimately, environmental gains.

Data center efficiency management. New approaches for improving data center efficiency through improved workload management are offered. A proof of concept for adaptive workload management by combining virtualization and statistical learning machines is published by researchers at the UC Berkeley Lab. The researchers claim a nearly 80 percent power savings in a simple workload test. (Bodik, Armbrust, Canini, Fox, Jordan & Patterson, 2008).

Alternative methods for modeling and managing data center airflow and cooling are also being explored. By conducting a proof of concept for real-time management of the thermal environment within the data center, IBM research fellows showed that data center air flow and cooling could be optimized by combining advanced dynamic metrics with measurement driven modeling techniques (Hamann, Kessel, Lyengar, Chung, Hirt, Schappert et al., 2009). The CFD thermal models currently utilized are static, generally time consuming to build, and are just starting to be embraced by the industry. The researchers believe that real time management of the data center thermal environment may provide a more practical alternative to these CFD models (Hamann, Kessel, Lyengar, Chung, Hirt, Schappert et al., 2009). The chip manufacturer AMD proposes that cooling power consumption could be made variable. Systems that adjust power consumption based on performance demands, together with intelligent cooling systems design, could adjust cooling consumption as well (AMD, 2005).

Alternative data center infrastructures are also explored. Enterprise data centers take approximately two years to build and have a life expectancy of 10-12 years. Therefore, data center capacity that is not fully utilized for 10-12 years is typically procured up front (Kaplan, Forrest & Kindler, 2008). There is a shift to modular data center infrastructure architecture, with each module sized to existing energy needs enabling steady growth (Loeffler, 2009).

Best Practices

Metrics. Currently there is no single set of universally accepted metrics for data center efficiency (Emerson, 2008). However, there is consensus in the industry that there is an urgent need to define and adopt metrics for IT performance that is scalable to the data center (Hamann, Kessel, Lyengar, Chung, Hirt, Schappert, et al., 2009). It does not have to be perfect. The basic criteria should be: will it drive the right behavior, will it be available at the device level, and is it scalable to the data center level (Emerson, 2008). Based on the latest move by the USEPA to use the PUE metric in its Energy Star compliance guidelines, the inclusion of PUE and DCiE metrics in the DOD's energy efficiency pamphlet, and the sheer number of vendors already utilizing these two metrics, The Green Grid's proposed metrics for Power Usage Effectiveness and Data center Infrastructure Efficiency are here to stay. No single framework or methodology for continuous data center efficiency management has gained significant traction as of yet.

Data center facility design. Within a given data center, the following practices are recommended: (1) server consolidation through virtualization and application stacking, (b) storage efficiency gain through deduplication, (c) power consumption forecasting, and (d) capacity management (Kaplan, Forrest & Kindler, 2008; Kumar, 2009; Long, Freeman & Patterson, 2009; Microsoft, 2008). There is clear agreement on the need for some basic data center design changes: minimizing air paths through hot/cold aisle layouts or moving CRAC to the rack, raising ambient temperature and chilled water temperatures, switching to variable speed motors for HVAC and IT equipment, and eliminating low-voltage transformers downstream from the UPS systems (Hewlett-Packard, 2008). General consensus on a set of IT equipment utilization practices also appears evident, with consolidation of data centers as the most obvious recommendation.

Potential Next Steps

As noted in the selected literature reviewed for this study, there are several recommended first steps that enterprise data centers can take when investigating efficiencies and power consumptions. Not surprisingly, these steps align with the basic business strategy: plan, implement, measure, and correct (Sokol, 2007).

- Diagnose – gather efficiency metrics. Enterprise data centers know how much is spend overall on energy, but they generally do not know where it is being consumed (Mainline, 2009). To manage consumption effectively it must first be measured. Application of the Green Grid metrics is a good start.

- Strategize – plan and design a path to implement best practices. This involves examination of innovative technologies and solutions to meet energy efficiency goals.

- Virtualize – a key strategy for server deployment. Enterprise data centers are most likely already employing virtualization to some degree; if so, this should be maximized thoughtfully. Do not join the one third of the corporations already drowning is virtual server sprawl (ESA, 2009).

- Manage – apply metrics to optimize efficiencies. Decrease idle server consumption & improve air management. Monitor infrastructure efficiency, use trending for improved planning, and employ heat and capping analysis. Adopt existing change control processes to include efficiency goals.

- Cool – focus on heating and cooling efficiency. Adopt latest cooling technologies and innovations such as water-cooling rack door mounts. Consider placing heavy computational workloads in data center locations that are easier to cool. Determine minimum effective ambient temperature.

References

ADC Inc. (2006). *TIA-942 Data center standards overview* [White paper]. ADC Telecom Inc. Retrieved from http://www.adc.com/us/en/Library/Literature/102264AE.pdf

ASE. (2009). *New study shows a staggering 4.7 million servers globally are doing nothing useful, wasting $25bn a year.* Washington, DC: Alliance to save energy. Retrieved from http://ase.org/content/news/detail/6190

AMD. (2005). *Power and cooling in the datacenter* [White paper]. Advanced Microsystems Devices. Retrieved from http://enterprise.amd.com/Downloads/34146D_PC_WP.pdf

AMD. (n.d.). *Power consumption and cooling in the data center: A Survey*. Advanced Microsystems Devices. Retrieved from http://enterprise.amd.com/Downloads/Ziff_Power_and_Cooling_IT_.pdf

Avelar, V., Torell, W., & Hampel, T. (2008). *Data center baseline study report* [White paper]. Beaverton, OR: The Green Grid. Retrieved from http://www.thegreengrid.org/~/media/WhitePapers/TGG_White_Paper_8_Data_Center_Baseline_Study_Report.ashx?lang=en

Battles, B., Belleville, C., Grabau, S., & Maurier, J. (2007). *Reducing datacenter power consumption through efficient storage* [White paper]. Sunnyvale, CA: Network Appliance, Inc. Retrieved from http://partners.netapp.com/go/techontap/matl/reducing-datacenter-power.pdf

Baroudi, C., Hill, J., Reinhold, A., & Senxian, J. (2009). *Green IT for dummies.* Hoboken, NJ: Wiley Publishing Inc..

Bell, C., & Smith, T. (2009). Critical evaluation of information sources. *University of Oregon Libraries.* Retrieved from https://blackboard.uoregon.edu/webapps/blackboard/content/contentWrapper.jsp?content_id=_3828671_1&displayName=Readings&course_id=_275381_1&navItem=content&href=http%3A//libweb.uoregon.edu/guides/findarticles/credibility.html

Best Practices. (n.d). In Business Dictionary Online. Retrieved from http://www.businessdictionary.com/

Bodik, P., Armbrust, M., Canini, K., Fox, A., Jordan, M., & Patterson, D. (2008). *A case for adaptive datacenters to conserve energy and improve reliability.* University of California at Berkeley. Retrieved from http://www.eecs.berkeley.edu/Pubs/TechRpts/2008/EECS-2008-127.pdf

Bouley, B. & Brey, T. (2009). *Fundamentals of data center power and cooling efficiency zones* [White paper]. Beaverton, OR: The Green Grid. Retrieved from http://www.thegreengrid.org/en/Global/Content/white-papers/Fundamentals-of-Data-Center-Power-and-Cooling

Brill, K. (2006). *Data center energy efficiency and productivity* [White paper]. Sante Fe, NM : The Uptime Institute. Retrieved from http://www.uptimeinstitute.org/symp_pdf/(TUI3004C)DataCenterEnergyEfficiency.pdf

Busch, De Maret, Flynn, Kellum, Le, & Meyers. (2005). Conceptual Analysis. Colorado State University Writing Guide. Retrieved from http://writing.colostate.edu/guides/research/content/com2b1.cfm

Carbon Footprint. (n.d.). In *Merriam-Webster's online dictionary*. Retrieved from http://www.merriam-webster.com

Carbon Footprint Ltd. (2009). *What is a carbon footprint?* Retrieved from http://www.carbonfootprint.com/carbonfootprint.html

Chen, P. (1976). *The entity-relationship model – toward a unified view of data*. Sloan School of Management, Massachusetts Institute of Technology. Retrieved from http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.123.1085&rep=rep1&type=pdf

Chillers. (n.d.). Retrieved Jan 12, 2009, from BettterBricks: http://www.betterbricks.com/DetailPage.aspx?ID=539

Data Center. (n.d.). Retrieved June 2, 2009, from Wikipedia: http://en.wikipedia.org/wiki/Data_center

Definition. (n.d.). In *Merriam-Webster's online dictionary*. Retrieved from http://www.merriam-webster.com

Dogra, R. (2010). Demand for data center storage to rise [Online exclusive]. *The Datacenter Journal.* Retrieved January 7, 2010, from http://datacenterjournal.com/content/view/3462/40/

Dunlap, K., & Rasmussen, N. (2006). The advantages of row and rack oriented cooling architectures for data centers [White paper]. West Kingston, RI: American Power Conversion retrieved from http://www.apcmedia.com/salestools/VAVR-6J5VYJ_R0_EN.pdf

Ebbers, M., Galea, A., Schaefer, M., & Khiem, M. (2009). *The green data center: An idea for today.* Armonk. NY : IBM. Retrieved from http://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/redp4523.html?OpenDocument

Emerson. (2008). *Energy logic: Calculating and prioritizing your data center IT efficiency actions* [White paper]. Columbus, OH: Emerson Network Power. http://www.emerson.com/edc/docs/EnergyLogicMetricPaper.pdf

Fan, X., Weber, W., & Barroso, L. (2007). *Power provisioning for a warehouse-sized computer* [White paper]. Mountain View CA: Google Inc. Retrieved from http://labs.google.com/papers/power_provisioning.pdf

Freeman, L. (2009). *Reducing data center power consumption through efficient storage* [White paper]. Sunnyvale, CA: NettApp, Inc. Retrieved from http://media.netapp.com/documents/wp-reducing-datacenter-power-consumption.pdf

Gartner, Inc. (April 2007). Gartner estimates ICT industry accounts for 2 percent of global C02 emissions. *Gartner Inc*. Retrieved from http://www.gartner.com/it/page.jsp?id=503867

Gartner, Inc. (October 2007). Gartner says data centers account for 23 per cent of global ICT CO2 emissions. *Gartner, Inc.* Retrieved from Gartner Newsroom: http://www.gartner.com/it/page.jsp?id=530912

Gartner, Inc (2009). *IT spending 2010.* Stamford, CT: Gartner Corporate Marketing.

Gonella, S., Lee, J., Chen, S., & Tian, C. (2009). *A practitioner's tool for carbon management in data centers*. New York, NY: IBM. Retrieved from http://domino.research.ibm.com/library/cyberdig.nsf/papers/B6EAC04C92D4B94D852575990053376E/$File/rc24774.pdf

Greenberg, S., Mills, E., Tschudi, B., Rumsey, P., & Myatt, B. (2006). *Best practices for data centers; Lessons learned from Benchmarking 22 data centers.* Proceedings of the 2006 ACEEE summer study on energy efficiency in Buildings. Retrieved from http://evanmills.lbl.gov/pubs/pdf/aceee-datacenters.pdf

Haas, J., Monroe, M., Pflueger, J., Pouchet, J., Snelling, P., Rawson, A., et al. (2009). *Proxy proposals for measuring data center productivity* [White paper]. Beaverton, OR: The Green Grid. Retrieved from http://www.thegreengrid.org/en/Global/Content/white-papers/Proxy-Proposals-for-Measuring-Data-Center-Efficiency

Hamann, H., Kessel, T., Lyengar, M., Chung, J., Hirt, W., Schappert, M., et al. (2009). Uncovering energy-efficiency opportunities in data centers. *IBM Journal of Research and Development. 53*(3),10.1-10.12.

Hewlett-Packard. (2008). *Energy efficiency in the data center* [White paper]. Hewlett-Packard Development Company. Retrieved from http://www.itwhitepapers.com/index.php?option=com_categoryreport&task=thankyou&title=4489&pathway=no&vid=1158

Kaplan, J., Forrest, W., & Kindler, N. (2008). *Revolutionizing data center energy efficiency* [White paper]. McKinsey & Company. Retrieved from http://www.mckinsey.com/clientservice/bto/pointofview/pdf/Revolutionizing_Data_Center_Efficiency.pdf

Koomey, J. (2007). *Estimating total power consumption by servers in the U.S. and the World.* Oakland, CA: Lawrence Berkeley National laboratory, Analytics press. Retrieved from http://enterprise.amd.com/Downloads/svrpwrusecompletefinal.pdf

Kumar, R. (2009). *Critical issues facing data center managers* (Report No. G00172282). http://www.gartner.com/DisplayDocument?id=1217721

Leedy, P. , & Ormrod, J. (2005). *Practical research: Planning and design*. Upper Saddle River, NJ: Pearson Prentice Hall.

Literature Reviews. (n.d.) The Writing Center. University of North Carolina at Chapel Hill. Retrieved from http://www.unc.edu/depts/wcweb/handouts/literature_review.html

Loeffler, C. (2009). Right-sizing your power infrastructure [White paper]. Cleveland, OH: Eaton Corporation. Retrieved from http://www.eaton.com/EatonCom/OurCompany/WhitePapers/idcplg?IdcService=GET_FILE&RevisionSelectionMethod=LatestReleased&Rendition=Primary&dDocName=WP09-21

Long, B., Freeman, J., & Patterson, M. (2009) *Assessment of EPA mid-tier data center at Potomac Yard* [White paper]. Beaverton, OR: The Green Grid. Retrieved from http://thegreengrid.com/en/Global/Content/white-papers/Assessment-of-EPA-Mid-Tier-Data-Center-At-Potomac-Yard

Loper,J. & Parr, S. (2007). *Energy efficiency in data centers: a new policy frontier*. Washington, DC: Alliance to Save Energy. Retrieved from http://www.fypower.org/pdf/ASE_DataCenter_EE.pdf

Mainline. (2009). *5 steps to a successful green computing solution* [White paper]. Tallahassee, FL: Mainline. Retrieved from http://whitepapers.techrepublic.com.com/abstract.aspx?docid=1104485

Mathew, P., Greenberg, P., Ganguly, S., Sartor, D., & Tschudi, P. (2009). *How does your data center measure up? Energy efficiency metrics and benchmarks for data center infrastructure systems.* Lawrence Berkeley National Laboratory. Retrieved from http://www.escholarship.org/uc/item/1dg6h3h2

Microsoft. (2008). *Best practices for energy efficiency in Microsoft data center operations* [White paper]. Redmond, WA: Microsoft Inc. Retrieved from http://download.microsoft.com/download/a/7/b/a7b72ab1-ca17-4589-923a-83b0ff57be6d/Energy-Efficiency-Best-Practices-in-Microsoft-Data-Center-Operations-CeBIT.doc

Mongan-Rallis, H. (2006). *Guidelines for writing a literature review*. University of Minnesota Duluth. Retrieved from http://www.d.umn.edu/~hrallis/guides/researching/litreview.html

Moore, J., Chase, J., Farkas, P., & Ranganathan, P. (2005). *Data center workload monitoring, analysis, and emulation*. Duke University. Retrieved from http://issg.cs.duke.edu/publications/caecw05-1.pdf

Nexsan. (2009). *The power/density paradox*. [White paper]. Thousand Oaks, CA : Nexsan Technologies. Retrieved from http://www.itwhitepapers.com/index.php?option=com_categoryreport&task=thankyou&title=6898&pathway=no&vid=1434

Obenzinger, H. (2005). *What can a literature review do for me*? Stanford University. Retrieved from http://ual.stanford.edu/pdf/uar_literaturereviewhandout.pdf

O'Flynn, A. (2008). Approach to green IT business at Fujitsu Australia Ltd. [Electronic version] *Fujitsu Science Tech Journal,45(1)*. Retrieved from http://web.cecs.pdx.edu/~karavan/cs533/papers/HPPAC-1569092686-paper-1.pdf

Ohara, D. (2008). *Sustainable computing- Putting IT on a diet.* Redmond, WA: Microsoft Inc. Retrieved from http://technet.microsoft.com/en-us/magazine/2008.05.sustainablecomputing.aspx?pr=blog

PG & E. (2006). *High performance data centers: A design guidelines sourcebook* [White paper]. San Francisco, CA: Pacific Gas and Electric Company. Retrieved from http://hightech.lbl.gov/documents/data_centers/06_datacenters-pge.pdf

Princeton. (2009). WordNet search: trend analysis. *Princeton University*. Retrieved from http://wordnetweb.princeton.edu/perl/webwn?s=trend%20analysis

Radicke, J., Roden, B., & Yunke, L. (2009). *Exploring views on data center power consumption and server utilization*. Lund University.

Rasmussen, N. (2007). *Electrical efficiency measurement for data centers* [White paper]. American Power Conversion. Retrieved from http://www.itwhitepapers.com/whitepaper5624/

Rawson, A., Pfleuger, J., & Cader, T. (2008). *Green Grid data center power efficiency metrics: PUE and DCIE* [White paper]. Retrieved from http://www.thegreengrid.org/en/Global/Content/white-papers/The-Green-Grid-Data-Center-Power-Efficiency-Metrics-PUE-and-DCiE

Routers. (2009). Company profile Gartner, Inc(IT) (NYSE Arca). Tompson Routers Corporate. Retrieved from http://www.reuters.com/finance/stocks/companyProfile?rpc=66&symbol=IT

Ruth, S. (2009). Green IT – More than a three percent solution? *IEEE Internet Computing, 13*(4), 74-78. Retrieved from http://www.computer.org/portal/web/csdl/doi/10.1109/MIC.2009.82

Seaton, I. (2004). *Best practices & concepts for computer room cooling* [White paper]. Westlake Village, CA: Chatworth Products,Inc. Retrieved from http://www.gocsc.com/uploads/white_papers/AFC24C0F06964C5897744B7BAF4C6363.pdf

Snevely, R. (2002). *Enterprise data center design and methodology*. Palo Alto, CA: Sun Microsystems. Retrieved from http://mediacast.sun.com/users/Kemer/media/bp_snevely_enterprise_data_center

Sokol, D. (2007). *Pleased but not satisfied*. David L. Sokol. self published hardcopy available from: http://www.amazon.com/Pleased-But-Satisfied-D-L-Sokol/dp/B0017OFXOQ

Stanley, J., Brill, K., & Koomey, J. (2006). *Four metrics define data center greenness* [White paper]. Santa Fe, NM: Uptime Institute. Retrieved from http://uptimeinstitute.org/wp_pdf/(TUI3009F)FourMetricsDefineDataCenter.pdf

SVLG Intel. (2009). *Control of computer room air conditioning using sensors in the IT equipment.* San Jose, CA: Silicon Valley Leadership Group. Retrieved from http://dcee.svlg.org/images/stories/pdf/Intel_cs.pdf

SVLG Power Assure. (2009). *Dynamic power management: Adjusting data center capacity in Real-Time.* San Jose, CA: Silicon Valley Leadership Group. Retrieved from http://dcee.svlg.org/images/stories/pdf/PowerAssure_cs.pdf

Symantec. (2009). *Green IT Global results: Survey results May 2009*. Mountain View, CA: Symantec Corporation. Retrieved from http://www.symantec.com/about/news/resources/press_kits/detail.jsp?pkid=greenitreport

TechTarget. (2009). *Computer room air conditioning units*[White paper]. Bitpipe.com. Retrieved from, http://www.bitpipe.com/tlist/Computer-Room-Air-Conditioning-Units.html

The Green Grid. (2009). The Green Grid releases the first comprehensive guide to European data centre legislation[press release]. Beverton, OR: The Green Grid. Retrieved Oct 18, 2009 from http://www.thegreengrid.org/~/media/press%20releases/TGG%20EMEA%20Policy%20whitepaper%20release_Cleaned%20up%20FINAL%20approved.ashx

Torres, J., Carrera, D., Hogan, K., Gavalda, R., Betran, V., & Poggi, N. (2008). *Reducing wasted resources to help achieve green data centers*. Technical University of Catalonia. Retrieved from http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=4536219

Turner, W., Seader, H., & Brill, K. (2005). *Industry standard tier classifications define site infrastructure performance* [White paper]. Santa Fe, NM: The Uptime Institute. Retrieved from http://www.dntp.com/news/pdfs/Data%20Center%20Tier_Classification.pdf

United States Department of Energy. (2008). *A quick start to energy efficiency* [Brochure]. Retrieved from http://www1.eere.energy.gov/femp/pdfs/data_center_qsguide.pdf

United States Department of Energy. (2008). *Greenhouse gases, climate change, and energy* [Borchure]. Retrieved from http://www.eia.doe.gov/bookshelf/brochures/greenhouse/greenhouse.pdf

United States Department of Energy. (2009). *The smart grid: An introduction* [Brochure]. Retrieved from http://www.oe.energy.gov/DocumentsandMedia/DOE_SG_Book_Single_Pages.pdf

United States Environmental Protection Agency. (2007). *Report to congress on server and datacenter energy efficiency* (Public Law 109-431).

http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf

United States Environmental Protection Agency. (2009). Enterprise server and data center energy efficiency initiatives. *Environmental Protection Agency*. Retrieved from http://www.energystar.gov/index.cfm?c=prod_development.server_efficiency

United States Environmental Protection Agency. (April 2009). *Energy Star Version 1.0Program requirements for computer servers*. Retrieved from http://www.energystar.gov/ia/partners/product_specs/program_reqs/computer_server_prog_req.pdf

United States Environmental Protection Agency. (August 2009). *Methodology for incorporating source energy use*. Retrieved from http://www.energystar.gov/ia/business/evaluate_performance/site_source.pdf

United States Environmental Protection Agency. (September 2009). *A guide to energy-efficient heating and cooling* [Brochure]. Retrieved from http://www.energystar.gov/ia/partners/publications/pubdocs/HeatingCoolingGuide%20FINAL_9-4-09.pdf

United States Environmental Protection Agency. (November 2009). *Energy Star Data center infrastructure rating development update* [pdf slides]. Web conference November 12, 2009. Environmental Protection Agency. Retrieved from http://www.energystar.gov/ia/partners/prod_development/downloads/Data_Center_Rating_Development_Results_Nov12.pdf

United States Government. (2008). Energy in brief. Retrieved from http://tonto.eia.doe.gov/energy_in_brief/greenhouse_gas.cfm

Wiedmann, T. & Minx, J. (2008). A definition of carbon footprint. In C. Pertsova(ED.), *Ecological Economics Research Trends* (pp. 1-11). Hauppauge NY: Nova Science Publishers. Retrieved