

Energy Efficient Data Centers

Prepared by

William Tschudi, Tengfang Xu, Dale Sartor, Bruce Nordman, Jon Koomey, and
Osman Sezgen

Environmental Energy Technologies Division
Lawrence Berkeley National Laboratory
Berkeley CA 94720

Final Report to

California Energy Commission

March 30, 2004

LBNL No. 54163

Legal Notice

Disclaimer

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor the Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Acknowledgements

Special thanks are extended to the Industry Partners and project participants for their support of this project.

- **Principal Investigator:**
 - Bill Tschudi, Lawrence Berkeley National Laboratory
- **Additional Investigators:**
 - Dale Sartor
 - Tengfang Xu
 - Bruce Nordman

Data Center Case Studies

- **Sub-contractors**
 - Peter Rumsey, Rumsey Engineers, Inc.
- **Industry Partners**
 - Network Appliance
 - Bank of America
 - State of California Franchise Tax Board
 - Equinix
- **Pacific Gas and Electric Company**
- **New York State Energy Research and Development Authority**
- **US Department of Energy Federal Energy Management Program**

Load Characterization in California

- **Additional Investigators:**
 - Jon Koomey
 - Osman Sezgen
 - John Steinmetz

Data Center Energy Research Roadmap

- **Subcontractor/Contributing Author:**
 - E-Source, Jay Stein
- **RMI Charrette:**
 - Houston Eubanks

Acknowledgements (Continued)

- **Numerous Data Center professionals**
- **7 X 24 Exchange**
- **Uptime Institute**

Preface

The Public Interest Energy Research (PIER) Program supports public interest energy research and development that will help improve the quality of life in California by bringing environmentally safe, affordable and reliable energy services and products to the market place.

The PIER program, managed by the California Energy Commission (Energy Commission), annually awards up to \$62 million to conduct the most promising, public interest energy organizations, including individuals, businesses, utilities, and public or private research institutions.

PIER funding efforts are focused on the following six RD&D program areas:

- Buildings End-Use Energy Efficiency
- Energy-Related Environmental Research
- Environmentally-Preferred Advanced Generation
- Industrial/Agricultural/Water End-Use Energy Efficiency
- Energy Systems Integration

What follows is the final report for the Energy Efficient Data Centers contract, contract number 500-01-024 conducted by the Lawrence Berkeley National Laboratory (LBNL). The report is entitled Energy Efficient Data Centers. This project contributes to the Industrial/Agricultural/Water End-Use Energy Efficiency program.

For more information on the PIER Program, please visit the Energy Commission Website <http://www.energy.ca.gov/pier/reports.html> or contact the Energy Commission at (916) 654 - 4628

Table of Contents

Executive Summary.....	1
Abstract	5
1.0 Introduction	7
1.1. Objectives	9
1.2. Report Organization	10
2.0 Project Approach.....	11
2.1. Data Center Electrical Load Characterization	11
2.2. Data Center Market Characterization.....	14
2.3. High-performance Data Centers – a Research Roadmap	14
3.0 Project Outcomes	16
3.1. Data Center Electrical Load Characterization	16
3.2. California Data Center Load Characterization	20
3.3. High-performance Data Centers – a Research Roadmap	21
4.0 Conclusions and Recommendations	22
4.1. Data Center Electrical Load Characterization	22
4.1.1. Conclusions	22
4.2. California Data Center Load Characterization	23
4.2.1. Commercialization Potential	23
4.2.2. Recommendations	23
4.2.3. Benefits to California.....	23
4.3. High-performance Data Centers – a Research Roadmap	24
4.3.1. Conclusions	24
4.3.2. Commercialization Potential	24
4.3.3. Recommendations	24
4.3.4. Benefits to California.....	25
References	26
Appendices	31

List of Figures

Figure 1. LBNL Data Center Energy Efficiency Website	2
Figure 2. Representative Energy End Use From An Actual Case Study.....	13
Figure 3. Measured and Projected Computer Load Intensity	16
Figure 4. Cumulative Distribution Of Computer Power Densities (UPS Power)	17
Figure 5. HVAC Load As A Percentage Of Total Load.....	18
Figure 6. UPS Efficiency vs. Load Factor	19
Figure 7 High-Performance Data Centers – a Research Roadmap Report.....	21

Executive Summary

Prior to this study, very little public information concerning the true electrical power requirements for California's data centers was available. This situation created much confusion, as data center developers claimed the need for large amounts of high-quality power – up to 250 Watts/square ft. To meet these requests, utilities receiving requests for new or upgraded service would have had to significantly upgrade the electrical infrastructure and/or provide for additional generating capacity. Uncertainty over electrical demand of present day and future information technology (IT) equipment led designers and operators of data centers to provide for unrealistically high electrical and HVAC system capacities. The Information Technology “industry” continually evolves and the prevailing wisdom was that the energy intensity of computing equipment would continue to rise, causing concern for the ability to provide cooling.

Prior case studies and limited investigation suggested that it should be possible to significantly improve the energy efficiency of data centers. To assist the PIER Industrial Program in identifying and prioritizing possible research areas, LBNL performed case studies involving six data centers and collaborated with various industry experts familiar with data center design and operation to develop a research agenda.

Project Objectives

The objectives of this project included:

- Obtaining measured energy end use energy efficiency information in 4-6 data centers, assembling additional data from synergistic projects sponsored by others, and using the case studies to help identify areas of potential public interest research.
- Characterizing the data center “market” in California
- Developing a research “roadmap” to guide California's public interest research on energy efficiency in data center facilities.

Project Outcomes

The project successfully arranged for data center's to be studied and obtained additional benchmarks beyond those originally planned. In all, benchmarks for sixteen data centers were obtained of which six were developed as part of this project. A key finding in the benchmarking results was that the average data center energy intensity today is on the order of 50 Watts/square ft. (compared to utility requests for up to 250 Watts/square ft.). Large variations in energy efficiency were also observed in the data center systems we studied, suggesting that there is room for significant improvement in data center performance using currently available technologies (such as improved chilled water system design or use of efficient air handlers instead of specialized computer room air conditioners. Individual case study reports were prepared for each

of the six centers as well as for most of the other synergistic case studies. These reports included a number of energy efficiency recommendations.

The market characterization task was challenging in that no market data exists concerning the number, location, or size of many types of data centers. Alternate methods of estimating the square footage of data centers were needed to try to bound the problem. Using estimated floor areas in combination with average measured energy intensities, an estimate of the total electrical demand in California in 2003 of between 250 and 375 MW was determined.

For the energy research roadmap, considerable interface with data center industry suppliers, owners, designers, and other researchers occurred. They provided valuable input, which coupled with observations from the case studies, led to a comprehensive public interest roadmap document. The roadmap is available from the LBNL data center energy efficiency website. (See <http://datacenters.lbl.gov>)



Figure 1. LBNL Data Center Energy Efficiency Website

Conclusions

The following conclusions can be drawn from this project:

- Little energy benchmark data exists for data center facilities.
- Energy intensities today are not as great as the industry would lead one to believe; results ranged from 4 to 65 W/ft².
- Energy Efficiency opportunities given today's technology are numerous.

- As a result of IT equipment loads continually changing, special considerations allowing for the large potential variations in loads are necessary for data center infrastructure systems to be energy efficient.
- Data Center industry professionals often lack knowledge of energy efficiency opportunities, but are eager to find solutions.
- Reliability and availability are key concerns – to be embraced, energy efficiency must align with them.
- Large energy savings should be attainable through further research and development.

Recommendations

We recommend that the PIER Industrial program adopt the roadmap for energy efficiency public interest research in data center facilities. High priority activities, if pursued, could yield near term savings while contributing to a longer-term integrated strategy. Industry estimates indicate that over 50% energy savings over current practice is possible. With the large and growing number of data center facilities in California, a permanent reduction of electrical demand would prevent or postpone utility expansions, improve reliability, and provide bottom-line savings to every industry that relies on data centers in its business.

Benefits to California

The case studies performed during this project have provided the data center owners with a clear understanding of their energy use and are likely to spawn energy efficiency improvement projects at the facilities that were studied. But a broader group of data center professionals were also exposed to the results of the studies through workshops and industry events. This has created increased industry awareness and has started dialogue within data center owners and designers, which will be very beneficial.

The High Performance Data Centers research roadmap provides the PIER program with much needed understanding of how the data center industry views needed research and its priority. The PIER Industrial Program will be able to utilize the roadmap to plan a strategy to aggressively make improvements in this critical market sector. The roadmap will also facilitate collaborations with other energy research and industry organizations thereby leveraging public sector efforts in California.

Abstract

Data Center facilities, prevalent in many industries and institutions are essential to California's economy. Energy intensive data centers are crucial to California's industries, and many other institutions (such as universities) in the state, and they play an important role in the constantly evolving communications industry. To better understand the impact of the energy requirements and energy efficiency improvement potential in these facilities, the California Energy Commission's PIER Industrial Program initiated this project with two primary focus areas: First, to characterize current data center electricity use; and secondly, to develop a research "roadmap" defining and prioritizing possible future public interest research and deployment efforts that would improve energy efficiency.

Although there are many opinions concerning the energy intensity of data centers and the aggregate effect on California's electrical power systems, there is very little publicly available information. Through this project, actual energy consumption at its end use was measured in a number of data centers. This benchmark data was documented in case study reports, along with site-specific energy efficiency recommendations. Additionally, other data center energy benchmarks were obtained through synergistic projects, prior PG&E studies, and industry contacts. In total, energy benchmarks for sixteen data centers were obtained.

For this project, a broad definition of "data center" was adopted which included internet hosting, corporate, institutional, governmental, educational and other miscellaneous data centers. Typically these facilities require specialized infrastructure to provide high quality power and cooling for IT equipment. All of these data center types were considered in the development of an estimate of the total power consumption in California.

Finally, a research "roadmap" was developed through extensive participation with data center professionals, examination of case study findings, and participation in data center industry meetings and workshops. Industry partners enthusiastically provided valuable insight into current practice, and helped to identify areas where additional public interest research could lead to significant efficiency improvement. This helped to define and prioritize the research agenda. The interaction involved industry representatives with expertise in all aspects of data center facilities, including specialized facility infrastructure systems and computing equipment. In addition to the input obtained through industry workshops, LBNL's participation in a three-day, comprehensive design "charrette" hosted by the Rocky Mountain Institute (RMI) yielded a number of innovative ideas for future research.

The Data Center Case Studies, the California Load Characterization Report, and the Energy Research Roadmap were completed and are attached as appendices to this report. They are also available through the LBNL Data Center website, <http://datacenters.lbl.gov/> along with other reference information.

1.0 Introduction

Background

Data Centers have long been an important component of California's industries, research organizations, educational facilities, and government. They are prevalent in both public and private-sector buildings serving many growing sectors of California's economy. Years ago, Data Centers containing mainframe computers were known to be very energy intensive (compounded by the associated high demands for air conditioning) requiring specialized infrastructure, but the development of the World Wide Web, and the shift to smaller, multiple-unit servers continued to utilize much of the same computer room infrastructure. Unfortunately, the continuous evolution of computing equipment creates uncertainty as to the overall energy intensity within data centers and the resulting demand on California's electrical power infrastructure. In the late 1990's, coincident with California's perceived energy crises, Internet hosting facilities were claiming that their infrastructure needed to support up to 200 Watts/square ft. and electric utilities even received requests for new power amounting up to 250 Watts/square ft. In some cases, in order to satisfy these requests, new power generation and/or transmission and distribution infrastructure would have been required raising issues of cost as well as complications of air-quality and power plant siting in urban areas. From California's public interest point of view, it was becoming very important to understand the patterns of energy consumption in data centers, and to be able to understand trends that could influence electrical load changes in the future.

Although the economic recession, beginning in 2001 in California, temporarily slowed and in some cases reversed growth in this market, the Information Technology (IT) industry continues to evolve. There is growing concern that technological advances are producing greater processing in smaller, more energy intensive devices. This simultaneously provides increased processing capability while complicating (or potentially making it impossible) to cool the devices. Energy efficiency improvements in servers and other IT equipment have not kept pace with the expanded processing capability. This situation has lead to several interesting scenarios involving growth in processing capability, and the increasing heat density trends in data centers. Some predict ever-increasing heat densities that eventually would force a change to liquid cooling. Other scenarios suggest that processing capability is outstripping demand and that this will result in compaction¹ and consolidation resulting in a reduction of cooling demands because in total, less computing equipment is needed. The case studies performed during this project identified evidence of both of these scenarios.

Prior investigations and anecdotal evidence suggested that there was considerable discrepancy between the electrical demand predicted when developers or building owners planned their data centers, and the actual measured electrical consumption. It

¹ Large inefficient computers replaced with smaller computers having much greater computational capacity.

was felt that overstating the heat load resulting from IT equipment electrical loads combined with the use of inefficient or outdated cooling practices often resulted in over-sizing of HVAC systems and consequent inefficient operation.

Because some data center professionals believe that IT equipment's energy intensities will continue to increase, there also was a tendency to exaggerate the impact of these facilities on the electrical power grid within California. (See: <http://n4e.lbl.gov/>.) Electric utilities felt that requests for electrical service were unrealistic but had no hard evidence to suggest otherwise. To attempt to provide some insight into this situation, an estimate of the electrical power requirements for California's data centers was needed.

The PIER Industrial Program recognized that improving energy efficiency in data centers represented an attractive public interest research opportunity and wanted a plan developed that would identify various activities that could be undertaken. The plan (termed a "Research Roadmap") was to be developed with content and priority input provided by data center building professionals. Working with industry owners, designers, and operators, the "Roadmap" would guide energy efficiency research and its adoption into the marketplace. Recent slowdowns in the internet economy created a window of opportunity to begin developing solutions before the next cycle of explosive growth puts further strain on the California electricity grid.

Computer facilities have one thing in common - they are extremely energy intensive. The case studies performed by LBNL identified significant opportunities for energy efficiency improvement in these buildings through better application of existing technologies and development of new approaches. Numerous opportunities are apparent within the individual systems that support Data Center operations. Further investigation into the interface of building systems and computer arrangements (ie. servers, racks, mainframe computers, etc.) reveals even more savings opportunities. Additionally, within the IT equipment itself, additional opportunities exist such as placing components in "sleep mode", designing circuits that use less energy, more efficient processors, more efficient computer software, and others. All of these opportunities were explored during the development of the energy research roadmap.

The PIER Data Center work was also leveraged through other synergistic projects and collaborations at LBNL, including:

NYSERDA - Case studies and energy benchmarking for two data centers in New York, a paper on energy benchmarking in data centers for the ACEEE Industrial conference², and a workshop with the 7 X 24 Exchange Organization, NY Chapter³.

² See: Appendix XVIII, ACEEE 2003 paper #162, "Data Centers and Energy Use, Let's Look at the Data"

³ See: Appendix XVII

Pacific Gas and Electric Company – Three prior case studies and energy benchmarking, and co-hosting a workshop with the Bay Area chapter of the 7 X 24 Exchange Organization (See: <http://datacenters.lbl.gov>.)

Industrial Partners – Many industrial partners provided in-kind support by providing input to the Data Center Energy Research Roadmap, and participating in the energy benchmarking.

Industry Associations – Informal collaboration with the Uptime Institute, the 7 X 24 Exchange Organization, a newly formed ASHRAE committee (TC 9.9) establishing standards for data center cooling, the CEETherm (a collaboration between the University of Maryland and Georgia Institute of Technology), and the Silicon Valley Manufacturers Group.

US Department of Energy – Federal Energy Management Program (FEMP) – Case studies and benchmarking for two federal data centers.

Prior work sponsored by the **US Environmental Protection Agency** including a report by Jennifer Mitchell-Jackson entitled Energy Needs in an Internet economy: A Closer Look at Data Centers. (See: <http://enduse.lbl.gov/Info/datacenterreport.pdf>.)

Rocky Mountain Institute (RMI) [Various Sponsors] – “Low Power Data Centers: Integrated Design Charrette”, a collaboration of over 75 data center professionals and researchers.

1.1. Objectives

The objectives of this project were to advance the knowledge of energy use in data centers, estimate the electrical demand of these facilities in California, investigate energy efficiency opportunities, and develop a research agenda which could be adopted by the PIER program. The structure of the project consisted of two primary tasks:

- Data Center Load Characterization - which included the individual case studies and energy end-use benchmarking, as well as an estimate of data center market in California.
- Development of a Research Roadmap to identify and prioritize energy research, and deployment of new technologies and strategies to improve data center energy efficiency.

1.2. Report Organization

This report addresses each of the two major tasks. The Project Approach, Project Outcome, and Conclusions and Recommendations sections of the report each contain separate summaries of the respective tasks. A brief summary of the task activity is included, and the project deliverables are attached as appendices. The appendices generally provide greater detail of the task, the findings, and recommendations.

This report is organized as follows:

Section 1.0	Introduction
Section 2.0	Project Approach
Section 3.0	Project Outcomes
Section 4.0	Conclusions and Recommendations

There are 19 Appendices.

Appendix I. “Estimating Total Power Used By Data Centers in California”; Jonathan G. Koomey, Osman Sezgen, and Robert Steinmetz, 2003

Appendix II. “Data Center Energy Characterization Study, Facility 1”; PG&E Case Study; Rumsey Engineers, Inc., 2001

Appendix III. “Data Center Energy Characterization Study, Facility 2”; PG&E Case Study; Rumsey Engineers, Inc., 2001

Appendix IV. “Data Center Energy Characterization Study, Facility 3”; PG&E Case Study; Rumsey Engineers, Inc., 2001

Appendix V. “Data Center Energy Benchmarking Case Study, Facility 4”; Federal Data Center Facility, Rumsey Engineers, Inc.; LBNL - TengFang Xu, 2003

Appendix VI. “Data Center Energy Benchmarking Case Study, Facility 5”; Federal Data Center Facility, Rumsey Engineers, Inc.; LBNL - TengFang Xu, 2003

Appendix VII. “Data Center Energy Benchmarking Case Study, Facility 6”; PIER Project, Rumsey Engineers, Inc.; LBNL – William Tschudi, 2003

Appendix VIII. “Data Center Energy Benchmarking Case Study, Facility 7”; PIER Project, Rumsey Engineers, Inc.; LBNL – William Tschudi, 2003

Appendix IX. “Data Center Energy Benchmarking Case Study, Facility 8”; PIER Project, Rumsey Engineers, Inc.; LBNL – William Tschudi, 2003

Appendix X. “Data Center Energy Benchmarking Case Study, Facility 9”; PIER Project, Rumsey Engineers, Inc.; LBNL – William Tschudi, 2003

Appendix XI. “NY Data Center Energy Benchmarking and Case Study, Facility 10”; NYSERDA Case Study, Syska & Hennessy; LBNL – William Tschudi, 2003

Appendix XII. “NY Data Center Energy Benchmarking and Case Study, Facility 11”; NYSERDA Case Study, Syska & Hennessy; LBNL – William Tschudi, 2003

Appendix XIII - “High-Performance Data Centers – a Research Roadmap”, Lawrence Berkeley National Laboratory Report No. 53483

Appendix XIV - RMI Charrette brochure, attendee list, and report.

Appendix XV - 8-22-02 Workshop Agenda, presentation, and meeting summary

Appendix XVI - 10-16-02 Workshop meeting notice and presentation

Appendix XVII - 4-17-03 Meeting notice and presentation

Appendix XVIII – ACEEE 2003 paper #162, “Data Centers and Energy Use – Let’s Look at the Data”

Appendix XIX – Annotated Bibliography

2.0 Project Approach

2.1. Data Center Electrical Load Characterization

Case Studies

Under this project, case studies were performed for six data centers in four very different organizations. For each, an energy end use breakdown was determined through actual energy measurement and analysis (benchmarking). The limited energy benchmark data were leveraged through other case studies performed through similar synergistic projects and by a leading data center industry association, the Uptime Institute. Each of the case studies provided potential energy efficiency improvement observations specific to the site.

To begin our investigation of the data center market in California, a sampling of various data centers were studied to determine their current energy use and the opportunity for energy efficiency improvement. First, the population of data centers to be included needed to be defined. For this project, a broad definition of a datacenter was adopted. This included various types of computing environments characterized by the requirement for specialized cooling systems and other specialized infrastructure such as raised floors, power conditioners, uninterruptible power supplies, etc. Using this definition, a number of industries such as Internet service providers were included as were most other industries that rely on large computing centers in their businesses such as banks, healthcare, etc. – virtually any large company. In addition, this market includes educational, governmental, research, and other institutions.

A diverse group of data center facilities were ultimately selected for the case studies and benchmarking. The participants included a computer disc drive manufacturer, an Internet hosting facility, a bank, and a California government facility. Two of the facilities had multiple data centers, which enabled two additional centers to be included within the project’s budget. At each site, meetings were held with the facility staff to describe the energy monitoring that would be required, and to collect any existing data along with selected design information. A subcontractor, Rumsey Engineers, then

visited the sites to obtain energy end use data. The work at the facility generally occurred over a 3-5 day period depending upon constraints at the participating site. Where energy use information was available, such as direct readout from uninterruptible power supplies, it was recorded. In other situations, energy-monitoring equipment such as clamp on power meters was used to measure actual energy use. All energy use within the data center was accounted for resulting in a total energy end use breakdown such as shown in Figure 2.




Figure 2. Representative Energy End Use From An Actual Case Study

In addition, the energy efficiency of key systems was determined. This generally included the HVAC chilled water system efficiency, the uninterruptible power supply (UPS) efficiency, transformer losses, etc. For HVAC, the metrics that were used provide benchmarks of the efficiency of making and delivering chilled water in terms of kW/ton. To calculate this metric, flow measurements were taken (or in some cases design data were used). To investigate the relative effectiveness of the HVAC systems in the various data centers, the ratio of HVAC power use to the total power consumption was determined. It was thought that a lower ratio may indicate a more energy efficient HVAC system, however other factors may also influence this. Efficiency of key electrical equipment (UPS, transformers, etc.) was determined by measuring input and output of the device.

It is common practice to express building electrical power requirements in terms of watts/square ft. (W/sf). To benchmark the relative energy intensity of the various load components in a data center, several different, and often confusing, building areas are used. These may include: gross area of the entire building, the area of raised floor, the area under computer racks, etc. To enable comparison of IT equipment's energy intensity to other industry data, we adopted the definition of "electrically active" space as defined by the Uptime Institute. This definition excludes support areas, storage areas, and major walkways thereby considering the energy intensity in the area where IT equipment is operating (generally, but not always on raised floors). This area data was obtained from design information and assistance from the host site. Using this area, an energy intensity of the "electrically active space" is developed since this is the area of interest that is housing IT equipment. Infrastructure support can likewise be expressed in terms of the energy intensity within this area.

Once data collection was complete, a site-specific report was prepared and then reviewed with the participant. These reports summarize the data collected and provide generic and specific recommendations for possible energy efficiency improvements. Generally, the reports are anonymous, as requested by the participating sites. Once the participant agreed with the content, the report was finalized and posted on the LBNL website: <http://datacenters.lbl.gov>.

2.2. Data Center Market Characterization

Characterizing the broadly defined “data center market” in California proved to be a daunting task. There currently is no comprehensive source of information concerning floor area or electrical demand for these facilities. Investigations with industry suppliers, designers, utilities, and other public sources of information yielded sketchy and unreliable data. For some types of centers, such as internet hosting facilities, prior industry studies, including real estate market estimates, were available and provided insight as to the size of that component of the market prior to the “dot com bust”. For others, such as banks or educational facilities, little information exists. For these types of data centers, estimates were attempted through various means such as server shipments, or amount of raised floor sold, but ultimately these avenues were not useful. Eventually some estimates were determined through discussions with industry experts, or use of other parameters such as student count for educational institutions. These methods result in highly uncertain estimates. A report detailing the methodology for these estimates is attached as Appendix I. The report also provides a methodology for estimating data centers’ contribution to electrical load broken down by the major public utilities in California.

2.3. High-performance Data Centers – a Research Roadmap

The philosophy employed in developing the research roadmap was to identify features of data centers where energy efficiency improvements were likely to be attainable through public interest research. These features were identified in a number of ways including LBNL observations during case study development, industry input through workshops, a charrette, individual consultation, participation in data center conferences, and interaction with industry associations and public interest organizations.

Early in the development of the roadmap, a workshop was held at the ACEEE summer study in Asilomar, CA with the goal of developing a framework for the roadmap. The meeting was held to capitalize on the fact that a number of leading energy experts were assembled, but was also well attended by some key industry experts. The meeting had key industry participation representing the design community; data center owners, energy engineers and consultants, and public goods program managers. Representatives of the following organizations provided input during this workshop:

AT&T
Pacific Gas and Electric
California Energy Commission
Southern California Edison
NYSERDA
ACEEE
Sure Power Corporation
Liebert Corporation
EYP Mission Critical Facilities
E-Source
Loudcloud
En-wise

The first workshop provided guidance for the project and a forum to identify barriers to efficient operation, areas of needed research, and other related research underway. The workshop agenda, LBNL presentation, and a summary of the workshop input are attached as Appendix XV. Following the workshop, LBNL developed the topic descriptions and suggested research areas in a first draft of the roadmap.

During the development of the roadmap, LBNL sought additional input from data center building design professionals, data center facility operators, and firms that provide specialty equipment such as computer room air conditioners. Additionally, research ideas were generated through informal collaborations with organizations such as the Uptime Institute (www.upsite.com), Intel Developers Forum, the 7 X 24 Exchange Organization (<http://www.7x24exchange.org/>), and CEETherm (see: <http://www.me.gatech.edu/me/publicat/brochures/Mettl/Bro0302.htm>). LBNL presented some of the preliminary project findings at the October 16, 2002 meeting of the Bay Area 7 X 24 Exchange Organization. Additional input was received at this meeting. A copy of the LBNL presentation and an unedited summary of the input are attached as Appendix XVI.

Once the first draft of the roadmap was developed, it was distributed to industry advisors for comment. LBNL further developed the roadmap topics through participation in a three-day charrette hosted by the Rocky Mountain Institute (RMI) where approximately 90 data center experts evaluated the energy efficiency potential and suggested numerous areas where current practice could be improved as well as where additional research could be expected to produce additional dramatic improvement. The RMI charrette brochure, attendee list, and charrette report is attached as Appendix XIV. The RMI report of the charrette was in preparation as of the completion of this project.

Finally, to further confirm the roadmap topics and to prioritize them, a workshop hosted by PG&E and the Bay Area chapter of the 7 X 24 Exchange Organization was held at the Pacific Energy Center on April 17, 2003. Benchmarking results were presented, and case studies were discussed including a detailed review of one case study (facility 8) presented by Rumsey Engineers. The attendees were then asked to provide input

concerning the priority of research tasks identified in the roadmap. The presentations and unedited summary of the input is attached in Appendix XVII.

Following the workshop, the roadmap was finalized and submitted to the California Energy Commission. The final roadmap (LBNL report no. 53483) is attached as Appendix XIII, and is available on LBNL's website: <http://datacenters.lbl.gov>.

3.0 Project Outcomes

3.1 Data Center Electrical Load Characterization

Case Studies

Each of the anonymous case study reports is available through the LBNL website: <http://datacenters.lbl.gov> along with case study reports developed through synergistic projects. Attached with these reports, is a detailed, annotated reference list providing other useful information.

Summary comparison data for selected metrics was prepared and presented to various industry organizations such as the 7x24 Exchange Organization. The summary presentation material is attached as Appendix XVII. This included a comparison of energy intensity as measured during these case studies. See Figure 3 below.


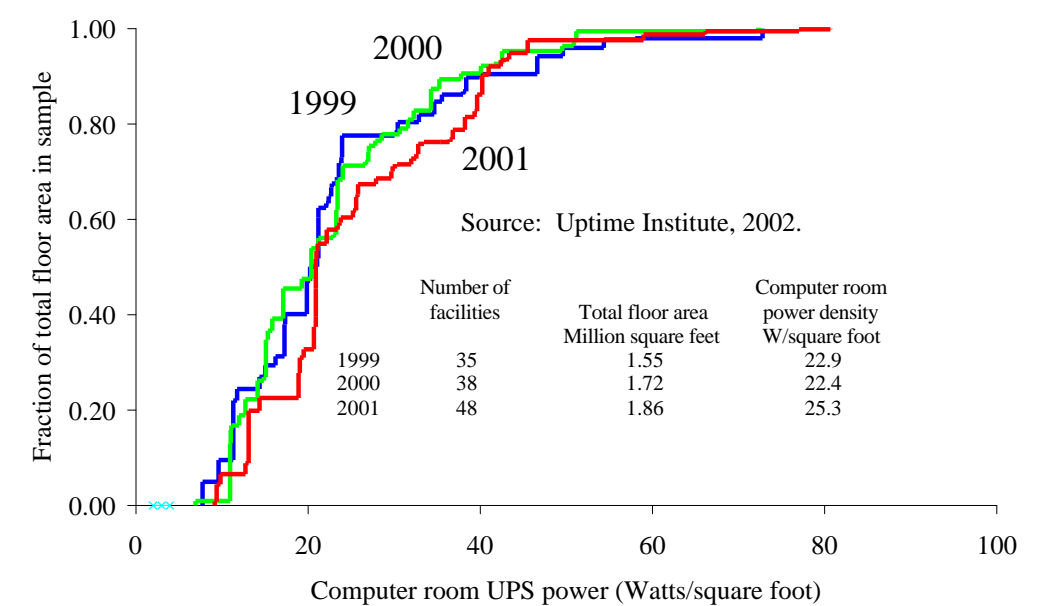


Figure 3. Measured and Projected Computer Load Intensity

Interestingly, these results demonstrated that the current average energy intensity of IT equipment in the measured data centers is approximately 25 Watts/SF. And, to project the intensity if the centers were full of similar equipment, the average intensity attributable to the IT equipment would only rise to approximately 40 Watts/SF. This

combined with the other infrastructure loads (HVAC, electrical losses, etc.) is much below the power needs that the industry was claiming in requests to utilities. To investigate how this compared to other industry benchmarks, the Uptime Institute was contacted and provided energy data that their member companies provided over a three-year period. This information was summarized and the relative distribution of the reported energy intensity is provided below in Figure 4.



Source: Uptime Institute, 2002.

Figure 4. Cumulative Distribution Of Computer Power Densities (UPS Power)

This information agrees well with measured results in the case studies, indicating that average energy intensity for computing equipment in 2001 averaged approximately 25 Watts/SF. The data also indicates that there was little change in overall intensity during the three years reported.

Insight into the relative effectiveness of HVAC systems is provided through the ratio of HVAC power to total power as shown in figure 4. The large variation in performance suggests that some of the HVAC system designs were significantly more efficient than others. For example, a system that had a 20% ratio was presumably providing cooling more efficiently than a system with a 50% ratio, although other factors, such as the amount of office space, or type/sizing of UPS systems, may enter into this ratio. To fully investigate whether the HVAC system was performing efficiently, other metrics such as kW/ton of chilled water should be examined. See Appendix XVII for more summary information.




Figure 5. HVAC Load As A Percentage Of Total Load

Surprisingly, the uninterruptible power supply (UPS) devices were found to be more inefficient than expected in many of the data centers based upon typical manufacturers claims. These devices are continuously consuming between approximately 5 to 50 % of the electrical power supplied to the IT equipment (compared to a nominal 10% that might be expected based upon manufacturer's data) with a multiplying effect approximately doubling the overall effect when considering the cooling of these loads.




Figure 6. UPS Efficiency vs. Load Factor

Figure 6 plots the actual measured UPS efficiency and as it illustrates, the efficiency of UPS systems drop off significantly at partial loads. This situation is prevalent in data centers that were studied. Data center UPS systems are operating at much less than full load for a variety of reasons including:

- Oversized equipment
- Partially filled with IT equipment
- Compaction due to replacement of old equipment with smaller, more efficient new equipment
- Redundancy strategies

3.2. California Data Center Load Characterization

The estimates of data center floor space in California are summarized below:

Hosting Facilities	2 million SF
Corporate Facilities	2-4 million SF
Institutional Facilities	0.5-1 million SF
Educational Facilities	0.5 million SF
Total Net Data Center Floor Area in California	5-7.5 million square feet

The average total energy intensity in California data centers (based upon case studies and Uptime Institute data) is approximately 50 Watts/square ft. Therefore, the total electrical demand attributable to California data centers is estimated at between 250 and 375 MW. See Appendix I for additional information.

3.3. High-performance Data Centers – a Research Roadmap



Figure 7 High-Performance Data Centers – a Research Roadmap Report

The report entitled “High-Performance Data Centers – a Research Roadmap” was developed through extensive interface with various industry experts in data center design and operation. The topics developed were validated and prioritized through several workshops and an intensive charrette. The completed roadmap is attached as Appendix XIII and is available for download from the LBNL data center website: <http://datacenters.lbl.gov/>

4.0 Conclusions and Recommendations

4.1. Data Center Electrical Load Characterization

4.1.1. Conclusions

Case Studies

Measured results for the data centers studied were quite revealing. Energy intensities were much lower than common industry claims. Average energy intensity for computing equipment was measured at approximately 25 W/SF with approximately equal energy intensity required for the infrastructure to support the IT equipment. This resulted in total energy intensity of approximately 50 W/SF. A wide variation in overall intensity and in individual system efficiencies was observed suggesting that significant improvement is possible with current best practices.

Data center owners/operators were eager to participate in the case studies. They wanted to know how their facility compared to others and were very open to energy efficiency recommendations. At each of the case study sites, the energy end use in the data center was determined along with the efficiency of key systems. Although many operators track the total IT equipment load – easily read from most UPS systems – very few had information on the total end use breakdown or the efficiency of their infrastructure systems. In addition, a number of efficiency recommendations were presented and most of the building owners indicated that some of the recommendations would be explored as retrofit projects, while others would be considered for future new construction. A follow-up to investigate what was adopted and the rationale for adoption could be the focus of a future PIER investigation.

Current practice using available technologies and techniques is far from optimal. By examining the better performing systems, current best practices may be able to be identified and should be the subject of a future examination. In addition, the case studies helped to identify areas where future research could lead to further efficiency gains. These areas were included in the research roadmap for future consideration.

Many in the industry hold a belief that the rising energy and computing intensity of IT equipment will lead to inability for it to be adequately cooled by air, yet the case studies and other industry data did not support this concern at this time for the overall data center. In fact in some data centers, compaction had taken place (large inefficient computers were replaced with smaller computers having much greater computational capacity - along with higher individual intensity). But the greater computational ability of the smaller, more energy intensive IT equipment resulted in lower overall energy intensities for the data center. While it is possible that the intensities may rise as the capability of these new machines are exceeded in the future, and additional equipment is added to existing computer rooms, there is a possibility that computational ability may continue to outpace needs - resulting in a net decrease in energy intensity.

4.2. California Data Center Load Characterization

Existing market data is not sufficient to accurately quantify the overall data center market in California. Based upon the various estimation methods in this study, the California data center electrical load is estimated to be 250-375 MW based upon estimates of floor area of 5 – 7.5 million SF. See Appendix I for additional details.

4.2.1. Commercialization Potential

Not Applicable

4.2.2. Recommendations

The limited number of case studies and benchmark results currently does not provide a statistically significant data set sufficient to conclusively bound the operating characteristics of California's entire data center population. While the results do suggest that there is significant possibility for improvement, additional benchmark results will be essential in order for a comprehensive best practices summary to be developed. It is recommended that additional benchmarks be obtained through PIER efforts, self-benchmarking, and other industry sources such as utility programs or industry associations such as the Uptime Institute. By determining the best performing systems in a large sampling of data centers, the "best practices" that led to their high performance can be identified. Some of the practices are likely to be applicable to both new construction and retrofit of existing data centers, depending upon their economic viability. These best practices could then be put into guidelines for data center owners and designers, forming a basis for public interest incentive programs.

As more benchmark data are obtained and best practices developed, additional deployment activities should be pursued to present this information to the target market. Continual involvement with industry associations such as the 7x24 Exchange, the Uptime Institute, Silicon Valley Manufacturers Group, etc. will help to ensure that energy efficiency has level of "visibility" on a par with current issues such as reliability, power quality, etc.

4.2.3. Benefits to California

The case studies performed during this project have provided the data center owners with a clear understanding of their energy use and are likely to spawn energy efficiency improvement projects at the facilities that were studied. But a broader group of data center professionals were also exposed to the results of the studies through workshops and industry events. This has created increased industry awareness and has started dialogue within data center owners and designers, which will be very beneficial.

4.3. High-performance Data Centers – a Research Roadmap

4.3.1. Conclusions

A consensus from numerous industry experts contacted during the roadmapping effort was that California's data centers need to become more energy efficient to save operating costs as energy costs rise, and to improve the reliability. The California Energy crisis and the downturn in “dot com” businesses have led data center owners to investigate efficiency opportunities. Companies contacted to participate in benchmarking and case studies were eager to participate in this project. Many were actively looking for efficiency opportunities. Some expressed the opinion that electrical distribution system reliability could be improved if many of the high intensity facilities could become more efficient. In addition increasing energy costs also emphasized the need for savings. As the data center market declined, the interest level for finding efficiency opportunities remained high because improving bottom line savings was essential. This, coupled with a perception that energy intensities were reaching the limits of air-cooling generated a lot of interest in researching new approaches. There is a unique opportunity at this time to couple energy efficiency with reliability concerns.

The roadmap focused primarily on energy efficiency aspects appropriate for public interest involvement. These activities developed through industry participation represent a large portion of the overall solution but are not all that can be done. The industry continues to research new technologies in very specialized areas where public interest research would not be appropriate.

4.3.2. Commercialization Potential

Although the roadmap addresses needed research - primarily those activities suited for public interest involvement - there are a number of topics where industry must take the lead, given encouragement or where a clear market potential exists. For example, more efficient heat transfer mediums within IT equipment would need to be developed by manufacturers, however public interest programs could provide needed incentives. IT professionals and data center owners must demand improvements in key equipment and building systems such as server power supply efficiency or uninterruptible power supply efficiency. In addition public interest programs such as ENERGY STAR labeling or utility incentives could be used to encourage market transformation.

4.3.3. Recommendations

The roadmap presents a multi-year research agenda. California's PIER program should proceed with the high priority tasks identified as the most beneficial to California companies by data center industry professionals. Collaboration with other industry efforts, such as the Uptime Institute, the 7 X 24 Exchange Organization, CEETHERM, IThERM, and ASHRAE, etc. should continue to enable as much of the roadmap to be realized as possible. To achieve the full potential in energy savings, progress on various levels is needed. Individual activities will achieve a level of improvement but attacking the overall opportunity will yield large benefits to industry and the state's electrical power industry's ability to provide adequate energy supply to meet demand. A multi-

year program that provides research into the full range of topics identified in the roadmap is expected to lead to a 40-50% overall reduction in energy use.

Information technology- its equipment - and the industries that depend upon it continually change. Computing technologies change rapidly and have a profound impact on the facilities in which they are located. Consequently, the roadmap topics and their priority should be reviewed periodically. The roadmap should be considered a living document and changes in priority and technological emphasis should be made as the market needs change.

4.3.4. Benefits to California

The High Performance Data Centers research roadmap provides the PIER program with much needed understanding of how the data center industry views needed research and its priority. The PIER Industrial Program will be able to utilize the roadmap to plan a strategy to aggressively make improvements in this critical market sector. The roadmap will also facilitate collaborations with other energy research and industry organizations thereby leveraging public sector efforts in California.

References

- ACEEE, and CECS. 2001. Funding prospectus for "Analysis of Data Centers and their implications for energy demand". Washington, DC, American Council for an Energy Efficient Economy (ACEEE); Center for Energy and Climate Solutions (CECS). July 2001.
- Aebischer, B., R. Frischknecht, C. Genoud, A. Huser, and F. Varone. 2002a. Energy- and Eco-Efficiency of Data Centres. A study commissioned by Département de l'intérieur, de l'agriculture et de l'environnement (DIAE) and Service cantonal de l'énergie (ScanE) of the Canton of Geneva, Geneva, November 15.
- Aebischer, B., R. Frischknecht, C. Genoud, and F. Varone. 2002b. Energy Efficiency Indicator for High Electric-Load Buildings. The Case of Data Centres. Proceedings of the IEECB 2002. 2nd International Conference on Improving Electricity Efficiency in Commercial Buildings. Nice, France.
- Anonymous. 2001. Model Data Center Energy Design Meeting. Austin Energy, Austin, TX, Feb 12-13.
http://www.austinenergy.com/business/energy_design_meeting.htm
- Anonymous. 2002a. 7 x 24 Update: Design & Construction - Issues and trends in mission critical infrastructure design, planning and maintenance.
<http://www.facilitiesnet.com/BOM/Jan02/jan02construction.shtml>. July 23, 2002. <http://www.7x24exchange.org/>.
- Anonymous. 2002b. Continuous Availability Review (CAR). The Uptime Institute: Computersite Engineering, Inc. <http://www.upsite.com/csepages/csecar.html>. July 22, 2002.
- Anonymous. 2002c. End-to-End Reliability Begins with the User's Definition of Success. The Uptime Institute.
<http://www.upsite.com/TUIpages/editorials/endtoend.html>. July 22, 2002.
- Anonymous. 2002d. Mechanical Systems Diagnostic Review (MSDR). The Uptime Institute: Computersite Engineering, Inc.
<http://www.upsite.com/csemsdr.html>. July 22, 2002.

- Anonymous. 2002e. Site Infrastructure Operations Review (SIOR). The Uptime Institute: Computersite Engineering, Inc. <http://www.upsite.com/csepages/cseior.html>. July 22, 2002.
- Baer, D. B. Emerging Cooling Requirements & Systems in Telecommunications Spaces, Liebert Corporation.
- Beck, F. 2001. Energy Smart Data Centers: Applying Energy Efficient Design And Technology To The Digital Information Sector. *Renewable Energy Policy Project (REPP): Washington, DC*. (November 2001 REPP).
- Blount, H. E., H. Naah, and E. S. Johnson. 2001. Data Center and Carrier Hotel Real Estate: Refuting the Overcapacity Myth. Lehman Brothers: TELECOMMUNICATIONS, New York, June 7, 2001. <http://www.lehman.com/>
- Bors, D. 2000. Data centers pose serious threat to energy supply. *Puget Sound Business Journal (Seattle)* - October 9, 2000. <http://seattle.bizjournals.com/seattle/stories/2000/10/09/focus5.html>
- Brown, E., R. N. Elliott, and A. Shipley. 2001. Overview of Data Centers and Their Implications for Energy Demand. Washington, DC, American Council for an Energy Efficient Economy, Center for Energy & climate Solutions (CECS). September 2001. <http://www.aceee.org/pdfs/datacenter.pdf.pdf> Callsen, T. P. 2000. The Art of Estimating Loads. *Data Center* (Issue 2000.04).
- Calwell, C., and T. Reeder. 2002. Power Supplies: A Hidden Opportunity for Energy Savings (An NRDC Report). Natural Resources Defense Council, San Francisco, CA, May 22, 2002. <http://www.nrdc.org/>
- Cratty, W., and W. Allen. 2001. Very High Availability (99.9999%) Combined Heat and Power for Mission Critical Applications. *Cinintel 2001*: 12. <http://www.surepowersystem.com>
- Elliott, N. 2001. Overview of Data Centers and their implications for energy demand. Washington, DC, American Council for an Energy Efficient Economy. Jan 2001, revised June 10, 2001.
- Feng, W., M. Warren, and E. Weigle. 2002. The Bladed Beowulf: A Cost-Effective Alternative to Traditional Beowulfs. *Cluster2002 Program* . <http://www-unix.mcs.anl.gov/cluster2002/schedule.html>; <http://public.lanl.gov/feng/Bladed-Beowulf.pdf>

- Frith, C. 2002. Internet Data Centers and the Infrastructure Require Environmental Design, Controls, and Monitoring. *Journal of the IEST* **45**(2002 Annual Edition): 45-52.
- Gilleskie, R. J. 2002. The Impact of Power Quality in the Telecommunications Industry. Palm Springs, CA, June 4. <http://www.energy2002.ee.doe.gov/Facilities.htm>
- Grahame, T., and D. Kathan. 2001. Internet Fuels Shocking Load Requests. *Electrical World* **Vol. 215** (3): 25-27.
http://www.platts.com/engineering/ew_back_issues.shtml
- Greenberg, D. 2001. Addendum to ER-01-15: A Primer on Harmonics. E-SOURCE, Boulder, Colorado, September 2001.
- Gross, P. 2002. Needed: New Metrics. *Energy User News*.
http://www.energyusernews.com/CDA/ArticleInformation/features/BNP_Features_Item/0,2584,82741,00.html
- Gruener, J. 2000. Building High-Performance Data Centers. *Dell Magazines - Dell Power Solutions* (Issue 3 "Building Your Internet Data Center").
http://www.dell.com/us/en/esg/topics/power_ps3q00_1_power.htm;
http://www.dell.com/us/en/esg/topics/power_ps3q00-giganet.htm
- Hellmann, M. 2002. Consultants Face Difficult New Questions in Evolving Data Center Design. *Energy User News*.
http://www.energyusernews.com/CDA/ArticleInformation/features/BNP_Features_Item/0,2584,70610,00.html
- Howe, B., A. Mansoor, and A. Maitra. 2001. Power Quality Guidelines for Energy Efficient Device Application - Guidebook for California Energy Commission (CEC). Final Report to B. Banerjee, California Energy Commission (CEC).
- Intel. 2002. Planning and Building a Data Center - Meeting the e-Business Challenge. Intel Corp.
http://www.intel.com/network/idc/doc_library/white_papers/data_center/. Aug 01, 2002.
- Koplin, E. 2000. Finding Holes In The Data Center Envelope. *Engineered Systems* (September 2000).

http://www.esmagazine.com/CDA/ArticleInformation/features/BNP_Features_Item/0,2503,8720,00.html

- Mandel, S. 2001. Rooms that consume - Internet hotels and other data centers inhale electricity. *Electric Perspectives* **Vol. 26** (No.3).
http://www.eei.org/magazine/editorial_content/nonav_stories/2001-05-01-rooms.htm
- Mitchell-Jackson, J. 2001. Energy Needs in an Internet Economy: A Closer Look at Data Centers, July, 2001.
- Mitchell-Jackson, J., J. G. Koomey, B. Nordman, and M. Blazek. 2001. Data Center Power Requirements: Measurements From Silicon Valley. *Energy—the International Journal* (Under review). <http://enduse.lbl.gov/Projects/InfoTech.html>
- Nordham, Reiss, and Stein. 2001. Delivering Energy Services to Internet Hotels and Other High Density Electronic Loads, Part I: Structure of the HiDEL Industry. Platts Research and Consulting, Boulder, CO.
- Patel, C. D., C. E. Bash, C. Belady, L. Stahl, and D. Sullivan. 2001. Computational Fluid Dynamics Modeling of High Compute Density Data Centers to Assure System Inlet Air Specifications. Reprinted from the proceedings of the Pacific Rim ASME International Electronic Packaging Technical Conference and Exhibition (IPACK 2001), © 2001, ASME.
- Patel, C. D., R. Sharma, C. E. Bash, and A. Beitelmal. 2002. Thermal Considerations in Cooling Large Scale High Compute Density Data Centers. 8th ITHERM Conference. San Diego CA.
- PG&E. 2001. Data Center Energy Characterization Study. Pacific Gas and Electric Company (subcontractor: Rumsey Engineers), San Francisco, Feb. 2001.
- Planet-TECH. 2002. Technical and Market Assessment for Premium Power in Haverhill. Planet-TECH Associates for The Massachusetts Technology Collaborative, <http://www.mtpc.org/>, Westborough, MA 01581-3340, Revision: February 20, 2002.
http://www.mtpc.org/NewsandReports/reports_other/Haverhill_Report.pdf;
<http://www.planet-tech.com/content.htm?cid=2445>
- RMI. 2003. Draft Final Report on Energy Efficient Data Centers - A Rocky Mountain Institute Design Charrette. Hayes Mansion Conference Center, San Jose, California.

- Robertson, C., and J. Romm. 2002. Data Centers, Power, and Pollution Prevention - Design for Business and Environmental Advantage. The Center for Energy and Climate Solutions; A Division of The Global Environment and Technology Foundation, June 2002. <http://www.cool-companies.org/> <http://www.getf.org>
- Roth, K. W., Fred Goldstein, and J. Kleinman. 2002. Energy consumption by office and telecommunications equipment in commercial buildings, Volume I: Energy Consumption Baseline. Arthur D. Little (ADL), Inc., 72895-00, Cambridge, MA, January 2002.
- Shields, H. and C. Weschler, 1998. Are Indoor Pollutants Threatening the Reliability of Your Electronic Equipment? Heating/Piping/Air Conditioning Magazine. May.
- Stein, Jay. 2002. More Efficient Technology Will Ease the Way for Future Data Centers. Proceedings 2002 ACEEE Summer Study on Energy Efficiency in Buildings.
- Sullivan, R. F. 2002. Alternating Cold and Hot Aisles Provides More Reliable Cooling for Server Farms. The Uptime Institute.
<http://www.upsite.com/TUIpages/whitepapers/tuaisles.html>
- The Uptime Institute. 2000. Heat-Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment. The Uptime Institute, Version 1.0., <http://www.upsite.com/>, <http://www.uptimeinstitute.org/heatdensity.html>
- Thompson, C. S. 2002. Integrated Data Center Design in the New Millennium. *Energy User News*.
http://www.energyusernews.com/CDA/ArticleInformation/features/BNP_Features_Item/0,2584,70578,00.html
- Uptime Institute, 2000. Heat Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment. Santa Fe, NM.
- Wood, L. 2002. Cutting Edge Server Farms - The blade server debate. [newarchitectmag.com](http://www.newarchitectmag.com).
- Uptime Institute, 2000. Heat Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment. Santa Fe, NM.

Appendices

Estimating Total Power Used By Data Centers in California:

Appendix I “Estimating Total Power Used By Data Centers in California”; Jonathan G. Koomey, Osman Sezgen, and Robert Steinmetz, 2003

Data Center Case Studies:

Appendix II. “Data Center Energy Characterization Study, Facility 1”; PG&E Case Study; Rumsey Engineers, Inc., 2001

Appendix III. “Data Center Energy Characterization Study, Facility 2”; PG&E Case Study, Rumsey Engineers, Inc., 2001

Appendix IV. “Data Center Energy Characterization Study, Facility 3”; PG&E Case Study, Rumsey Engineers, Inc., 2001

Appendix V. “Data Center Energy Benchmarking Case Study, Facility 4”; Federal Data Center Facility, Rumsey Engineers, Inc.; LBNL - TengFang Xu, 2003

Appendix VI. “Data Center Energy Benchmarking Case Study, Facility 5”; Federal Data Center Facility, Rumsey Engineers, Inc.; LBNL - TengFang Xu, 2003

Appendix VII. “Data Center Energy Benchmarking Case Study, Facility 6”; PIER Project, Rumsey Engineers, Inc.; LBNL – William Tschudi, 2003

Appendix VIII. “Data Center Energy Benchmarking Case Study, Facility 7”; PIER Project, Rumsey Engineers, Inc.; LBNL – William Tschudi, 2003

Appendix IX. “Data Center Energy Benchmarking Case Study, Facility 8”; PIER Project, Rumsey Engineers, Inc.; LBNL – William Tschudi, 2003

Appendix X. “Data Center Energy Benchmarking Case Study, Facility 9”; PIER Project, Rumsey Engineers, Inc.; LBNL – William Tschudi, 2003

Appendix XI. “NY Data Center Energy Benchmarking and Case Study, Facility 10”; NYSERDA Case Study, Syska & Hennessy; LBNL – William Tschudi, 2003

Appendix XII. “NY Data Center Energy Benchmarking and Case Study, Facility 11”; NYSERDA Case Study, Syska & Hennessy; LBNL – William Tschudi, 2003

High-Performance Data Centers – a Research Roadmap

Appendix XIII - “High-Performance Data Centers – a Research Roadmap”, Lawrence Berkeley National Laboratory Report No. 53483

Appendix XIV - RMI Charrette brochure, attendee list, and report.

Appendix XV - 8-22-02 Workshop Agenda, presentation, and meeting summary

Appendix XVI - 10-16-02 Workshop meeting notice and presentation

Appendix XVII - 4-17-03 Meeting notice and presentation

ACEEE Paper

Appendix XVIII – ACEEE 2003 paper #162, “Data Centers and Energy Use – Let’s Look at the Data”

Appendix XIX – Annotated Bibliography

Appendix I “Estimating Total Power Used By Data Centers in California”; Jonathan G. Koomey, Osman Sezgen, and Robert Steinmetz, 2003

ESTIMATING TOTAL POWER USED BY DATA CENTERS IN CALIFORNIA

Jonathan G. Koomey, Osman Sezgen, and Robert Steinmetz

INTRODUCTION

There are big uncertainties in estimating data center floor areas and associated electric loads. Four broad categories of data centers are of interest here:

- (1) Multi-tenant hosting facilities, which include data centers owned by third parties that house servers owned by other companies (co-location), as well as data centers that sell server services to companies who do not want to manage their own servers (managed hosting). These facilities are often housed in buildings devoted solely to these activities, or in buildings associated with telecommunications and networking equipment.
- (2) Corporate or enterprise data centers, owned by corporations and managed in-house. These data centers are often housed within existing facilities and may comprise only a small portion of the total floor area associated with those facilities.
- (3) Institutional and government data centers—these facilities are owned and operated by Federal, state, and local governments or by non-profit institutions.
- (4) Educational data centers serving students and faculty in post-secondary institutions.

This study estimates total power use for such facilities, focusing mainly on California. We begin by summarizing what is known about total floor area for these facilities, and then analyzing the power densities associated with this floor area. Total power use is then the total floor area times the power densities.

Unfortunately, there is little definitive work on aggregate floor areas of data center facilities in California (or elsewhere, for that matter). This study relies on back of the envelope calculations and on the data that do exist to assess the order of magnitude of this floor area. More research will surely be needed to fully understand the current and future floor stock of data centers in California—this report is only a first cut.

DEFINITIONS

Data center power requirements are also often misstated because of the lack of common metrics and terminology (Mitchell-Jackson et al. 2002, Mitchell-Jackson et al. 2003). Power densities for these facilities are often given in terms of watts per square foot (W/ft^2) or watts per square meter (W/m^2) but it is not always clear what such numbers mean. For example, a stated power density of $100 \text{ W}/\text{ft}^2$ could refer to the power drawn by an isolated rack of computer equipment or the average power density of the building.

Extrapolating the power requirements of a data center from an isolated rack is misleading because the floor area within a building includes office space, restrooms and hallways, all of which require much less power per square foot than does a computer rack.

Many estimates of total building floor area for data center facilities include significant amounts of space not critical to a data center's main function. In a data center facility, 15% to 75% of the building's area is usually designated for meeting rooms, offices, restrooms and hallways, but only the computer room floor area is relevant to assessments of the growth in electricity demand of data centers. Computer room floor area includes both the area under the racks of computer equipment and that of some common areas, such as the aisles between racks of computer equipment. It does not include office space, lobbies, bathrooms, space set aside for potential future expansion, or mechanical equipment rooms.

Because of the confusion over this issue, in our earlier work (Mitchell-Jackson et al. 2002, Mitchell-Jackson et al. 2003) we defined two key terms: *computer power density* is the power drawn by the computer equipment (in watts) divided by the computer room floor area (in square meters or square feet); *total computer room power density* is defined as the power drawn by the computer equipment and all of the supporting equipment such as power distribution units, uninterruptible power supplies, HVAC and lights (in watts) divided by the electrically active ("net") computer room floor area (in square meters or m²). Total computer room power density is the most meaningful indicator of power needs, but computer power density can also be useful, as we shall see below. Both parameters can be compared between buildings of different sizes as well as between data centers at different stages of development.

Some in the industry (e.g. Ken Brill of the Uptime Institute) are beginning to focus on power/rack as a more convenient and accurate metric. This choice eliminates the confusion over area-based metrics.

ESTIMATING FLOOR AREA

Most of the data that exist relate to hosting facilities, because that is where the growth was during the Internet boom, and that is what consultants found profitable to track. While fairly accurate information already exists for total multi-tenant web hosting space, corporate, institutional, and educational data center floor areas are not currently tracked by any industry sources.¹ Because of this lack of published information, we contacted sources in the relevant sectors to get their institutional perspectives (their contact information is contained in Table 1). In all cases we focus on the so-called "net floor area" or "electrically active floor area" of the computer rooms in data center facilities.

¹ At the PG&E/7X24 Exchange workshop On April 17, 2003, some participants expressed opinions about the amount of data center floor area in California. Email was sent to those participants to obtain whatever data existed, but no participants supplied useful data. This experience reflects the sorry state of information on data center floor area—lots of people have opinions on the topic, but there's little real data.

Multi-tenant hosting facilities

Two comprehensive studies for data-center floor area for hosting facilities were published in August 2000 by Salomon Smith Barney (Mahedy et al. 2000) and in January 2001 by Robertson Stephens (Juarez et al. 2001). Each study surveyed about 60 companies and developed U.S. and global distributions of hosting-type data centers by number and by floor area. They also estimated floor-area growth for the near future. In April 2001 Salomon Smith Barney (Mahedy et al. 2001) revised their floor area estimations, reducing them on a global basis by about 25% (unfortunately the later SSB report did not split out California or US data centers, so more geographically detailed correction factors are not available). Table 2 summarizes the estimates and projections published by these two companies. There is good agreement between these sources—especially after the update of the SSB numbers in 2001-- on the global gross floor area and the split of the floor area between managed hosting and colocation.

Looking into the situation in California in more detail, SSB 2000 (Mahedy et al. 2000) and RS (Juarez et al. 2001) tracked the companies presented in Table 3. Table 3 also shows the number of locations of each company in the three major areas in California, namely, San Francisco/Bay Area, Los Angeles/San Diego, and Sacramento. RS tracked more locations compared to SSB (about twice as many). However, for the locations tracked by SSB 2000, very detailed information is available, including estimates of net and gross floor area.

The floor area data for the California SSB locations are presented in Table 4. We focus in particular on the end of year 2000 data, because that is the same year used by the RS report. First, the estimate of net floor area in California (1.14 Msf) needs to be adjusted downward by 25% to reflect the global correction estimated in the SSB 2001 report (this correction reflects the slowdown in the high tech industry that first began in early 2000). Next, the SSB data must be adjusted upwards by a factor of 2 to reflect the larger coverage of the RS report. These sequential calculations yield an estimate of $1.14 \times 0.75 \times 2$ or 1.7 Msf as a lower bound estimate of the net floor area of hosting-type facilities in California. This calculation assumes that the high tech slowdown hit all regions equally, and that the additional coverage of the RS report over the SSB 2000 report is the same everywhere, and does not differ in California.

We know that even the RS report did not survey every relevant company in the hosting industry, but we expect that they captured the majority. This implies that roughly 2 Msf is a reasonable estimate for the net floor area of hosting-type facilities in California (with the understanding that this estimate is highly uncertain).

Corporate data centers

The uncertainty in total floor area is even greater for corporate than for hosting facilities, because details on corporate facilities are normally regarded as closely-held proprietary information. For this reason, no reliable data exist on corporate data center floor area in the aggregate. We rely instead on expert judgments to bound the range.

Our first attempt to estimate corporate data center floor area used total server shipments, estimates of equipment lifetimes, and typical floor areas associated with different types of equipment. Our second attempt involved estimating typical facility sizes by industry and estimating the percentage of firms of a given type that have in-house data centers. Time and budget constraints prevented us from conducting an exhaustive survey, so we assumed reasonable ranges by industry type and firm size and combined these estimates with California Census data to estimate total floor space. The projections generated by these two approaches were highly sensitive to the assumptions and as a result the range of possible projections was too large for us to have confidence in them. We therefore abandoned these approaches.

After being foiled in these initial attempts, we talked to several experts on this topic, most notably Ken Brill of the Uptime Institute and Ion Yadigaroglu of Koch Industries. We also asked some of the largest companies in the industry to supply us with their California floor areas, some of whom replied, some of whom didn't. After these discussions, we adopted a range of from 2 Msf to 4 Msf for corporate data center net floor area in California. This estimate is also highly uncertain.

Institutional and government data centers

Institutional facilities include those for non-profit organizations (including some hospitals). Government data centers include facilities at the Federal, state, county, and city levels. Unfortunately, we were only able to locate credible data at the state government level. Federal space operated by such organizations as the U.S. Air Force, DOE, FBI, CDC, and several others is difficult to inventory due to lack of centralized information and heightened security concerns.

State government data center space is largely concentrated within three large facilities all with active floor space in excess of 40,000 Sq. Ft. as well as several second tier data centers. According to State data center managers, larger cities within California each run their own data center, as well as larger counties. As an example, the County of Marin currently operates a 1850 sq. ft. data center for county related government purposes.

Table 5 shows the state government floor area estimates, which total almost 0.2 million square feet. Unfortunately no other compilations exist for the institutional and governmental data centers associated with other organizations within California, and the effort needed to collect these data was beyond the scope of this study. For purposes of this study, we estimate a range of 0.5 Msf to 1 Msf for all institutional and governmental data centers, including state government facilities, but these estimates cry out for further research and verification.

Educational facilities

To assess the floor area in educational facilities, a sample was taken of four public and private universities of varying sizes. We contacted staff at Stanford University, the University of California at Berkeley, the University of San Diego, and San Francisco State. We assumed that high schools and other secondary educational institutions would

not typically house data centers. The active data center space totals were then combined with total figures for full time student enrollment to arrive at an average data center floor area per student of approximately 0.5 square feet per student. This figure was then used to estimate total floor area based on the total California University and College Student population of approximately 1 Million. Table 6 depicts this calculation, which results in an estimate of about 0.5 million square feet of data center floor area in the education sector.

Summary of data center floor area in California

With hosting facilities at 2 Msf, corporate facilities between 2 and 4 Msf, institutional data centers between 0.5 and 1 Msf, and educational data centers at 0.5 Msf, our estimates yield a range for total data center net floor area in California of between 5 and 7.5 Msf.

ESTIMATING COMPUTER ROOM POWER DENSITIES

Multi-tenant hosting facilities

To estimate total computer room power density for these facilities, we reviewed billing data for five hosting type data center facilities throughout the country (Mitchell-Jackson 2001, Mitchell-Jackson et al. 2002). Electricity billing data were used to find average demand in the month with the highest consumption (this month was usually the most recent one). The highest average power demand for the facility was then divided by the computer room floor area. This estimate of total computer room power density is an overestimate because it assumes that all of the power for the entire facility is used for the computer room; thus our estimate is an upper limit for the total computer room power density. Even these overestimates, however, indicate that the total computer room power density was always less than 430 W/m² (40 W/ft²).² For comparison, typical office buildings have peak electrical demands of 54 to 86 W/m² (5 to 8 W/ft²).

Total computer room power density is partly a function of the occupancy level of the space and of the racks within the space. The average rack in the Bay area data center that we studied most closely was only one-third full (the average for 14 case studies covering various types of data centers was 60% full). Unfortunately, we don't have similar occupancy data for the other data centers analyzed in the Mitchell-Jackson et al. studies, but our experiences in visiting other data centers and in talking with people who design such facilities leads us to conclude that the occupancy level for that Bay area data center of about one-third full was probably typical for hosting-type facilities.

The six hosting facilities benchmarked in the main part of this report span a range of power densities that is closer to that shown in the Uptime institute's data set on power densities for corporate facilities than to the facilities described earlier in this section of

² The four values were: 8 W/ft², 34 W/ft², 35 W/ft², and 38 W/ft².

this appendix. It is not known which data are most representative for all U.S. hosting facilities

Corporate data centers

For corporate data centers we summarize new data that we obtained recently from the Uptime Institute in Santa Fe, New Mexico <<http://www.upsite.com/>>. To the best of our knowledge, these (along with the four corporate data centers examined in the main part of this report) are the first available and reliable measured data on power used by corporate data centers.

The data are from quarterly surveys for facilities that were members of the Site Uptime Network, for the years 1999, 2000, and 2001. In 2001, the year with the largest sample size, the number of facilities in the sample is 48, representing about 1.9 million sq. feet of net (electrically active) floor space in total. While this sample is not necessarily a statistically representative one, it does include data centers of many different types with a wide range of computer room power densities.

Figure 1 summarizes the results for all three years. Note that these computer room power densities do not include electricity used for cooling and auxiliary equipment, an omission for which we correct below. Mean computer power densities are between 22 and 25 W/ft² over this three-year period, with the maximum being about 80 W/ ft² and the minimum less than 10 W/ft². It is not possible to discern time trends with these data given the small number of data points and the short time period.

Ken Brill of the Uptime institute estimates that total loads in data centers are typically 1.6 to 2.5 times the computer room UPS loads and our case studies confirm similar ranges of performance. For simplicity, we adopt a multiplier of 2, which yields the result that total computer room power density for these facilities ranges from 44 to 50 W/ ft². This total is slightly higher than that for hosting facilities, which we expect because corporate facility managers often have more control over the planning and operation of their facilities than do those managing hosting facilities.

Institutional and government and educational data centers

With the exception of the four institutional/government data centers described in the case study reports, we were unable to locate publicly available and representative data on power use for the broader population of institutional, government, and educational data centers, although this situation will change as the benchmarking activities in California and other states begins to bear fruit (Tschudi et al. 2003). The four institutional/government data centers in the main part of this report span a range that is comparable to that shown in the Uptime data for corporate facilities, though we do not know if these four facilities are representative.

Summary of power density estimates

For purposes of the calculations in this appendix, we adopt an estimate of total computer room power density of 540 W/m² (50 W/ft²) for hosting, corporate, institutional, and

educational data centers (based on the detailed data for corporate facilities). It is not known with precision how the power densities of corporate facilities compares to those of hosting, institutional, governmental, and educational data centers, though the range from the benchmarked hosting and institutional facilities reported in the case study reports is comparable to the ranges in the Uptime institute data set for corporate data center power densities.

ESTIMATING TOTAL DATA CENTER LOADS IN CALIFORNIA

A range of 5 to 7.5 Msf of floor area and an average total computer room power density of 50 W/sf yields estimated total loads for data center facilities in California in 2003 of between 250 and 375 MW. Assuming a load factor of 0.9, these loads imply electricity consumption of between 2 and 3 TWh/year.

One of the questions of concern to California policy makers is how the data center loads are distributed by utility. We address this question with the rough calculation shown in Table 7. This table distributes the electricity use estimated in the previous paragraph to major California utilities based on each utility's share of non-residential electricity use in 2001. Unfortunately, data limitations prevented us from tallying actual data center loads by utility, and true bottom-up electricity use data may differ substantially from those shown here. This initial rough calculation is the best that anyone can do until Tier 1 research completes their data collection and analysis in the coming two years (see below).

RECOMMENDATIONS FOR FUTURE WORK

The case studies for NY and California, combined with planned additional benchmarking activities, may yield insights that will help us better understand these loads. Power densities in the buildings under study may allow us to refine the power densities used above. Similarly, contacts in the industry developed from these case studies may also lead us to better estimates of raised floor area in California and the U.S. as a whole, which is an even more pressing need.

We have been frustrated in our efforts to collect credible data on floor area of corporate and institutional data centers. These facilities are often quite small and difficult to track—details on their characteristics are also often treated as trade secrets. We tried many different industry sources, including the main supplier of raised floor in the US, to no avail. There are many opinions in the industry, but none of these sources had any hard data. We also attempted two scoping calculations using business type, server shipments, rack occupancy, rack placement density, and equipment lifetimes, but these calculations were not sufficiently constrained to yield useful results. There were many possible combinations of plausible assumptions, and they led to a wide range of estimates that didn't help us much.

Currently there is no single research source that is tracking comprehensive data center floor space and capacity. Estimates of total capacity, even on a regional basis are rough extrapolations based upon small samples and industry experience. However, Tier 1

Research, a Minneapolis based market research firm recognized as the best source in data center statistics and tracking, along with AFCOM, the leading data center managers association, has recently begun an effort to resolve this issue. According to Tier 1 Research President Andrew Schroepfer, their approach is to build a comprehensive database of facilities nationwide including enterprise facilities. They expect to complete this work within the next 12 to 18 months (i.e. by the end of calendar year 2004). Tier 1 Research currently tracks multi-tenant hosting space and keeps estimates of capacity that it regularly updates via a list of 883 hosting facilities worldwide.

CONCLUSIONS

It is surprisingly difficult to obtain market data on the floor area associated with data centers (particularly corporate and institutional facilities). We continue to explore methods for estimating this number. We believe that the power density estimates are more reliable than those for floor area, because they are based on actual data from about 2 Msf of raised floor area from across the country.

Our initial estimate is that total data center loads in California are between 250 and 375 MW. The uncertainty in these estimates is very large, It is also important to note that the concern over data center loads was largely focused on the hosting facilities, which were predominantly (though not exclusively) new loads, while most of the corporate, institutional, and educational data centers have existed for decades. Only 40% of the lower power use estimate and a quarter of the higher power use estimate are associated with hosting facilities in California, and the rest are associated with corporate, institutional, government, and educational data centers. There is also great uncertainty in the future floor areas and loads, as the data center industry is evolving rapidly.

REFERENCES

Blount, Harry E., Henry Naah, and Ethan S. Johnson. 2001. *Data Center and Carrier Hotel Real Estate: Refuting the Overcapacity Myth*. New York, NY: Lehman Brothers. June 7.

Juarez, Richard A. 2000. *Virtual Bricks II: Virtual Econ 101 Update: A Comprehensive Guide for Understanding eCommerce Infrastructure Evolution and Convergence*. Boston, MA: FleetBoston Robertson Stephens Inc. May. Available from World Wide Web: (<http://www.rscs.com>)

Juarez, Richard A., Michael T. Alic, Karkhaniz Chetan S., and Brett D. Johnson. 2001. *Space Dex III—Hosting Space: Not All Space Is Created Equal—Smart, Complex Space Takes Center Stage*. Boston, MA: Robertson Stephens Inc. January 29. Available from World Wide Web: (<http://www.rscs.com>)

Mahedy, Stephen, Dan Cummins, and Danny Joe. 2000. *Internet Data Centers: If Built... Will They Come?* New York, NY: Salomon Smith Barney. August 3.

Mahedy, Stephen, Dan Cummins, and Danny Joe. 2001. *Internet—Infrastructure and Service Providers*. New York, NY: Salomon Smith Barney. April 18.

Mitchell-Jackson, Jennifer. 2001. *Energy Needs in an Internet Economy: A Closer Look at Data Centers*. M.S. Thesis, Energy and Resources Group, University of California, Berkeley.

Mitchell-Jackson, Jennifer, Jonathan Koomey, Michele Blazek, and Bruce Nordman. 2002. "National and Regional Implications of Internet Data Center Growth." *Resources, Conservation, and Recycling (also LBNL-50534)*. vol. 36, no. 3. October. pp. 175-185.

Mitchell-Jackson, Jennifer, Jonathan Koomey, Bruce Nordman, and Michele Blazek. 2003. "Data Center Power Requirements: Measurements From Silicon Valley." *Energy—the International Journal (also LBNL-48554)*. vol. 28, no. 8. June. pp. 837 - 850.

Tschudi, William, Priya Sreedharan, Tengfang Xu, David Coup, and Paul Roggensack. 2003. *Data Centers and Energy Use – Let's Look at the Data*. Proceedings of the 2003 ACEEE Summer Study on Energy Efficiency in Industry. Rye Brook, NY: American Council for an Energy Efficient Economy, in Washington, DC. July.

Table 1: Contact information for sources of data center floor area in California

<i>Corporate/Enterprise</i>	
Ron Spangler Product Manager Liebert 614-888-0246	Bill Perry Technical Product Manager Tate Access Floors, Inc. 410-799-4200
Andy Schroepfer President Tier1 Research andy@tier1research.com 763-694-9992	Michael Dell Lydia Leong Senior Research Analysts Gartner, Inc. 408-468-8000
<i>Government</i>	
Al Smith Chief Technology Officer Teale Data Center 916-464-3934	Doug Grandy Chief Energy Policy Advisor State of California 916-375-4403
Cal Zissel IT Manager San Francisco International Airport 650-821-3350	Dennis Lionberger IT Director Marin County Civic Center 415-499-6314 DLionberger@co.marin.ca.us
<i>Education</i>	
Barabara Morgan Director of Strategic Technology Planning, ISNT University of California barbm@uclink.berkeley.edu 510-642-5567	John Vier Data Center Manager Stanford University 650-723-4004
Kevin Barney Data Center Facilities Manager UC Berkeley 510-642-5378	Jack Tse Director of IT Operations San Francisco State University 415-338-2627
Tom Cofflin Facilities Capital Planning Manager University of San Diego 619-921-8508	

Table 2: Floor Area for Internet Data Centers (million sq. ft.)

	Date	Estimation Year (end of year)	Coverage	Gross Floor Area	Net Floor Area	Data Center Type Co-location /Managed hosting	Number of Facilities
Salomon Smith Barney (Mahedy et al. 2000)	Aug. 2000	2000 mid-year	Global	15	8	70%/30%	
		2000		36	19		
		2001		60	33	67%/33%	
		2000 mid-year	US	8	5		
(Mahedy et al. 2001)	Apr. 2001	2000		18	10		
		2001		32	18		
		2000	Global (1)	27	14		
		2001		43	22		
Robertson Stephens (Juarez 2000)	May 2000	1999	Global	7			
		2000		21			
	Jan. 2001	2000	Global	25		69% / 31%	517
		2001		44		76% /24%	810
Lehman Brothers (Blount et al. 2001)	Jan. 2001	2001	North America	43		50-60% co-location	

Sources: Mahedy et al. 2001, Mahedy et al. 2000, Juarez et al. 2001, Juarez et al. 2000, Blount et al. 2001.

- (1) SSB 2001 is not as detailed as SSB 2000. It is not possible to separate domestic floor area from international.

Table 3a. California Sites for Managed Hosting

	Salomon Smith Barney (estimate for end-of year 2000)			Robertson Stephens (end-of-year 2000)		
Company	SF/Bay Area	LA/San Diego	Sacramento	SF/Bay Area	LA/San Diego	Sacramento
Applied Theory				2		
ConXion				1		
Digex				1		
Digital Island				1		
EDS						1
Exodus	5	1		7	3	
Futurelink					1	
Genuity	3	1		2	1	
IBM	1					
Intra				1		
Navisite	2			2		
Usinternetworking	1				1	
Total	12	2	0	17	6	1

Sources: Juarez et al. 2001 and Mahedy et al. 2000.

Table 3b. California Sites for Colocation Facilities

	Salomon Smith Barney (end-of year 2000)			Robertson Stephens (end-of-year 2000)		
Company	SF/Bay Area	LA/San Diego	Sacramento	SF/Bay Area	LA/San Diego	Sacramento
AboveNet				4	1	
AT&T Worldnet	1	2		2	1	
Cable&Wireless	1			1	1	
Colo.com	3	2		3	2	
Concentric	1			1	1	
Data Return Hosting	1					
Equinix				1		
First World	2	1		1	4	
Globix	2	1		1		
Hostcentric				1		
HostPro(Micron)		1			1	
iAsiaWorks				3		
Inflow		1			2	
Intel	1			2		
Level 3				3	2	
NetNation Communications					1	
PSINet		1			1	
Qwest Communications	3			3	1	
Switch & Data Facilities	1	2	1	1	3	1
TierraNet, Inc.					1	
Universal Access				2	1	
UUNet/Worldcom	1			2	1	1
Verio				5	1	
Wavve Telecom						2
Williams Communications				1	1	
Total	17	11	1	37	26	4

Sources: Juarez et al. 2001 and Mahedy et al. 2000.

Table 4. California Data Center Hosting Floor Areas (sq.ft.)

Company	Location	2000(Mid-year)		2000(Estimate)		2001(Estimate)	
		Gross	Net	Gross	Net	Gross	Net
AT&T	CA-LA			65,000	50,000	65,000	50,000
AT&T	CA-San Diego	7,475	5,750	19,500	15,000	19,500	15,000
AT&T	CA-San Francisco	7,475	5,750	19,500	15,000	19,500	15,000
Cable & Wireless	CA-LA (TBD)						
Cable & Wireless	CA-Santa Clara	85,800	66,000	85,800	66,000	85,800	66,000
Colo.com Hosting	CA-LA	34,710	17,355	34,710	17,355	34,710	17,355
Colo.com Hosting	CA-San Diego			22,068	11,034	22,068	11,034
Colo.com Hosting	CA-San Francisco	20,576	10,288	20,576	10,288	20,576	10,288
Colo.com Hosting	CA-San Ramon			18,677	9,339	18,677	9,339
Colo.com Hosting	CA-Santa Clara			25,000	12,500	25,000	12,500
Concentric-NETLINK	CA-San Jose	10,000	5,000	10,000	5,000	10,000	5,000
Data Return Hosting	CA-San Francisco	2,500	1,250	2,500	1,250	2,500	1,250
	CA-Sacramento						
EDS Hosting	(TBD)						
Exodus	CA-LA	123,000	61,500	123,000	61,500	123,000	61,500
Exodus	CA-Santa Clara	25,000	12,500	25,000	12,500	25,000	12,500
Exodus	CA-Santa Clara	52,000	26,000	52,000	26,000	52,000	26,000
Exodus	CA-Santa Clara	150,000	75,000	150,000	75,000	150,000	75,000
Exodus	CA-Santa Clara	125,000	62,500	125,000	62,500	125,000	62,500
Exodus	CA-Santa Clara			150,000	75,000	150,000	75,000
FirstWorld	CA-San Diego	20,000	13,340	20,000	13,340	20,000	13,340
FirstWorld	CA-San Francisco			40,000	26,680	40,000	26,680
FirstWorld	CA-Santa Clara	19,000	12,673	19,000	12,673	19,000	12,673
Genuity	CA-LA	10,900	5,450	10,900	5,450	10,900	5,450
	CA-Mountain						
Genuity	View	146,000	73,000	146,000	73,000	146,000	73,000
Genuity	CA-Palo Alto	10,900	5,450	10,900	5,450	10,900	5,450
Genuity	CA-San Jose	10,900	5,450	10,900	5,450	10,900	5,450

Source: Mahedy et al. 2000.

Table 4 (continued). California Data Center Hosting Floor Areas (sq.ft.)

Company	Location	2000 (Mid-year)		2000 (End-of-year)		2001 (End-of-year)	
		Gross	Net	Gross	Net	Gross	Net
Globix Hosting	CA-LA			100,000	50,000	100,000	50,000
Globix Hosting	CA-Santa Clara	62,000	23,000	62,000	23,000	62,000	23,000
Globix Hosting	CA-Santa Clara			100,000	50,000	100,000	50,000
HostPro	CA-LA	13,500	10,000	20,250	15,000	20,250	15,000
IBM	CA-LA					25,000	16,250
IBM	CA-San Jose			25,000	16,250	25,000	16,250
IBM	CA-San Jose					25,000	16,250
IBM	CA-Sunnyvale					25,000	16,250
Inflow	CA-San Diego	16,800	12,000	16,800	12,000	16,800	12,000
Intel	CA-Santa Clara	85,000	43,000	85,000	43,000	85,000	43,000
Navisite	CA-San Jose	33,000	22,000	37,950	25,300	43,643	29,095
	CA-Scotts						
Navisite	Valley	6,000	4,000	6,000	4,000	6,000	4,000
PSINet	CA-LA	5,000	25,000	5,000	25,000	5,000	25,000
	CA-San						
PSINet	Francisco (TBD)						
Qwest	CA-LA					100,000	65,000
	CA-San						
Qwest	Francisco	50,000	25,000	50,000	25,000	50,000	25,000
Qwest	CA-San Jose			100,000	65,000	100,000	65,000
Qwest	CA-San Jose					100,000	65,000
Qwest	CA-Sunnyvale	50,000	25,000	60,000	30,000	72,000	36,000
Qwest	CA-Sunnyvale					100,000	65,000
Switch and Data	CA-LA	35,800	17,900	35,800	17,900	35,800	17,900
	CA-Oakland						
Switch and Data	(TBD)						
Switch and Data	CA-Sacramento			27,200	13,600	27,200	13,600
Switch and Data	CA-San Diego	28,000	14,000	28,000	14,000	28,000	14,000
	CA-San						
Switch and Data	Francisco (TBD)						
Switch and Data	CA-San Jose			40,000	20,000	40,000	20,000
USInternetworking	CA-Milpitas	17,000	15,000	18,700	16,500	20,570	18,150
UUNet	CA-San Jose	20,000	10,000	20,000	10,000	20,000	10,000
	CA Total	1,283,336	710,156	2,043,731	1,142,859	2,438,294	1,398,054

Source: Mahedy et al. 2000.

Table 5. Floor area of state government data centers, by major agencies

	Floor area Total Sq. Ft.
Teale Data Center	43,100
Health and Human Services	45,000
Hawkins Data Center (Law Enforcement)	45,000
Franchise Tax Board	15,000
CALPERS	10,000
Legislature	10,000
California Highway Patrol	1,000
Department of Motor Vehicles	1,000
Employemnt Development Department	1,000
Secretary of State	1,000
Dept. of Corrections	1,000
Department of Veterans Affairs	1,000
Public Utilities Commission	7,000
Air Resources Board	
Health Services	
Social Services	
Industrial Relations	
Department of Forestry	
Department of General Services	
Dept. of Consumer Affairs	
Dept. of Mental Health	
State Treasurer	
State Controller	
California Youth Authority	
Board of Equalization	
Total California State Government Space:	181,100

(1) Source: Al Smith: Teale data center CTO for the State of California.

Table 6: Back of the envelope estimate of educational data center floor area

	Data Center Space (sf)	Full Time Students	Avg sf per student
Stanford	18,000	18,500	0.97
UC Berkeley	13,000	33,000	0.39
University of San Diego	3,500	6,000	0.58
San Francisco State	4,500	27,000	0.17
Total/Average	39,000	84,500	0.46
Extrapolated Total for all California Full Time University/College Students	473,878	1,026,736	0.46

Sources : California Department of Finance: California State University, University of California, California Postsecondary Education Commission: Community Colleges, Independent Colleges & Universities, Private 2-year Colleges, Other Public Institutions.

Table 7: Rough distribution of California data center electricity use by utility

Utility	2001 California non-residential electricity use by utility % of total	Low estimate of data center electricity use GWh	High estimate of data center electricity use GWh
PG&E	35%	700	1000
SCE	36%	700	1100
SDG&E	6%	100	200
LADWP	10%	200	300
Other	14%	300	400
Total	100%	2000	3000

(1) Source of non-residential electricity use fractions in 2001:

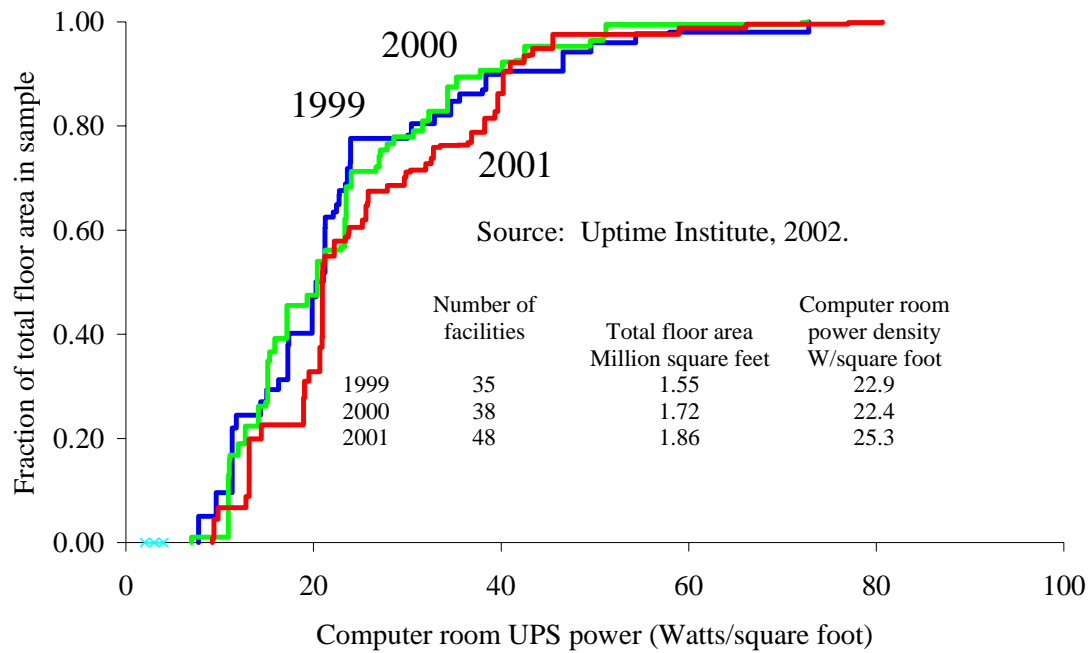
http://www.energy.ca.gov/electricity/utility_sales.html

Downloaded 18 June 2003 by Jonathan Koomey

(2) Total data center electricity use spread across utilities in proportion to the non-residential electricity use associated with each utility.

(3) "Other" utilities include SMUD and all other municipal and private utilities in California.

Figure 1: Cumulative distribution of computer power densities (UPS power)



Source: Uptime Institute, 2002.

**Appendix II. “Data Center Energy Characterization Study,
Facility 1”; PG&E Case Study; Rumsey Engineers, Inc., 2001**

**DATA CENTER ENERGY
CHARACTERIZATION STUDY
SITE REPORT**

DATA CENTER FACILITY 1

FEBRUARY 2001

I. Review of Site Characteristics

Facility

Facility 1 is a 102,500 square foot (sf) data center facility located in Silicon Valley, California. This facility provides co-location service, which is an unmanaged service that provides rack space and network connectivity via a high-capacity backbone. The building houses 62,870 sf of data center space, 7,310 sf of office space, 25,170 sf of support space and 5,390 sf of electrical room space. Roughly 60 percent of the building is data center space. Combined office and support space accounts for 32 percent of the building. 75 percent of the data center was occupied during the monitoring period. The data center's environmental system operates 24 hours a day year-round. The users of the co-location space require full access to and control of their caged space 24 hours a day.




Customer Cage

Electrical Equipment and Backup Power System

Electricity use at Data Center Facility 1 was monitored from December 19, 2000 through December 26, 2000. Facility 1 has a PG&E service feed of 4,000 kVA. Of this, an average of 3,540 kVA is being used. The site drew an average of 2,990 kW over the period of December 23, 2000 through December 26, 2000. The load varied $\pm 3\%$ from the average, and the load factor was 0.96.

Chart 1. Whole Building Power Consumption



The facility utilizes Liebert uninterruptible power supplies (UPS) to condition the utility power in order to shield critical loads on the data center floor from disturbances. In addition, a constant delivery voltage to the data center is maintained. Anytime the voltage drops below 480, the UPS systems' batteries feed in the necessary voltage to maintain 480 volts. The UPS converts AC current and stores it as DC current. When the voltage is needed, it is converted back to AC current. In the event of a power loss, 4 Cummins diesel generators each provide 2,000 kW each for backup. At fifty percent load (4,000 kW), the generators with a full tank of diesel fuel can provide electricity to the building for approximately 40 hours. The data center is designed for N+1 redundancy on the electricity circuit level.



UPS Module

Cooling System

The mechanical pad includes four Evapco 340 ton, closed circuit water-cooling towers and condenser water pumps. This cooling tower plant provides condensed water to the computer room air conditioning (CRAC) units, and to office space air conditioning units and heat pumps located throughout the building. Each cooling tower has two two-speed fans and a circulating water pump. The condensed water loop consists of four pumps, with three pumps operating at a time. The cooling tower plant is designed for a redundancy of 15 percent.

The data center is on a raised floor, through which cooling air is circulated via the CRAC units. Forty-eight CRAC units, with a cooling capacity of twenty-two tons each, serve the space. The room temperature and relative humidity are maintained at an average of 64 °F and 42 %, respectively. Because it was built out in phases, the data center is divided into three different areas.



Data Center CRAC Unit

Facility 1 has four electrical rooms, each equipped with four CRAC units with a cooling capacity of ten tons apiece. A temperature of 68 °F and relative humidity of 45% are maintained in the electrical rooms. Both types of cooling units have a compressor EER (energy efficiency rating) of 11.7.

II. Electricity Use Characteristics

Facility A's end-use of electricity is shown below in Table 1 and Chart 2. "Other" was calculated by subtracting all of the measured data from the "Whole Building" power. Over 75 percent of the power goes to energizing the data center: 51 percent for server and related equipment loads, and 25 percent to the cooling equipment. Furthermore, 34 percent of the power goes to the HVAC equipment. This is a significant amount of the whole building power consumption and is where energy efficiency opportunities can be taken. A 15 to 50 percent reduction in HVAC electricity use can be achieved. This corresponds to 152 to 506 kW of electricity savings. "Other" contributes 13 percent of the whole building consumption; it includes items such as losses of power in the electrical equipment, office plug loads, and domestic water heating.

The power energizing the data center floor remains relatively constant throughout the day. This can be verified by examining the Chart 4 which shows the power drawn by the CRAC units and server loads in a section of the data center over a period of three days. The server load power varied $\pm 0.5\%$ from the average. This indicates that power drawn remains the same regardless of the peak time usage of the servers and supporting equipment. The CRAC unit power varied $\pm 9\%$ from the average, demonstrating that the internal heat load is fairly constant and that the building shell load has little effect on the environment.

Table 1. End-Use of Electricity

Description	Electricity Consumption (kW)	% of Whole Building	Square Feet (sf)	Watts / sf
Whole Building	2,990	--	102,500	29
Data Center Server Load	1,500	51	62,870	24
Data Center CRAC Units	738	25		12
Cooling Tower Plant ¹	114	4		1.8
Electrical Room Cooling	118	4	5,390	22
Office Space Conditioning	41	1	7,310	5.6
Lighting	73	2	102,500	0.7
Other	402	13	102,500	3.9

¹ Assume that all cooling tower plant power is used for data center cooling. This is a simplifying assumption.

Chart 2. Facility 1 Electricity End-Use




Chart 3. Facility 1 Operating Energy Densities





Chart 4. Constant Power Drawn by Data Center CRAC Units and Server Loads




III. Electricity Use Diversity

In determining the size of the equipment needed in a data center facility, designers, in most cases, use an energy density value expressed in watts per square foot (W/sf). The type and number of server equipment in the data center are difficult to estimate; designers therefore oversize electrical equipment, so that a lack of capacity will not be a concern. Table 2 shows the operating conditions of the facility in comparison with the designed conditions. An extrapolated value was also calculated to determine what the operating W/sf would be if the data center were fully loaded.

Table 2. Server Load Diversity Factor

Measured W/sf	24
Extrapolated Full Load W/sf	32
Design W/sf	50
Diversity Factor (Measured / Design)	0.48
Diversity Factor (Extrapolated / Design)	0.64

Chart 5. Server Load Density



The data center was only 75 percent occupied at the time of monitoring. Thus 75 percent of the 62,870 sf area was used in calculating the extrapolated data center server load. The extrapolated diversity factor demonstrates that when the data center is fully occupied, it will operate at 64 percent of the expected or designed load.

**Appendix III. “Data Center Energy Characterization Study,
Facility 2”; PG&E Case Study, Rumsey Engineers, Inc., 2001**

**DATA CENTER ENERGY
CHARACTERIZATION STUDY
SITE REPORT**

DATA CENTER FACILITY 2

FEBRUARY 2001

I. Review of Site Characteristics

Facility

Facility 2 is an 118,700 square foot (sf) data center facility located in Silicon Valley, California. The facility consists of four levels and is located in a two-story building. The building houses 60,400 sf of data center space, 58,300 sf of office space, 6,800 sf of support space, and 4,600 sf of electrical room space. This facility provides co-location service, which is an unmanaged service that provides rack space and network connectivity via a high-capacity backbone. Data center space occupies approximately half of the building area. During the monitoring period, the data center was 65 percent occupied. Combined office and support space accounts for 55 percent of the building area. The data center's environmental system operates 24 hours a day year-round. The users of the co-location space require full access to and control of their caged space 24 hours a day.




Customer Cage


Electrical Equipment and Backup Power System

The electricity use at Data Center Facility 2 was monitored from January 3, 2001 through January 10, 2001. An average kVA of 4,190 and an average of 4,000 kW were being drawn over the monitoring period. The load factor over the period was 0.96, and the load varied $\pm 4\%$ from the average.

The facility utilizes five dynamic uninterruptible power supply (UPS) system modules. (See picture below). The generators condition the incoming utility power via the choke coil in order to remove any irregularities in the signal. This is achieved by having a constantly spinning generator in parallel with the main service feed. The induction coupling is also kept spinning so that the generator will continue producing electricity in the event of a power loss. In such an event, the Mitsubishi diesel engines will start up and engage to the induction coupling via the clutch to keep the generator turning. Each generator at full load can provide 1,500 kW apiece for 40 hours. The electrical system is designed for N+2 redundancy.

Chart 1. Whole Building Power Consumption





Uninterruptible Power Supply Module

Mechanical Systems



Condenser Units

The mechanical pad on the building rooftop consists of 75 air-cooled condensers for the CRAC (computer room air conditioning) units, five exhaust fans, one outside air supply fan, five package air conditioning/heating units, three humidifiers and two hot water boilers. The hot water boilers are for the four office package units. The rooftop mezzanine houses the UPS systems, nine exhaust fans, and four electric block heaters. The block heaters are used to pre-heat the diesel engine block of the UPS system.

The data center is on a raised floor, through which cooling air is circulated via the CRAC units. Seventy Data Aire CRAC units with a nominal capacity of twenty-five tons each, serve the data center space. In addition to the CRAC units, a 40 ton package unit adds humidification and cooling to the data center. The room temperature and relative humidity are maintained at an average of 70 °F and 45 %, respectively. The data center is divided into 11 rooms and each have a cooling system designed for N+1 redundancy. The EER (energy efficiency rating) of the Data Aire CRAC unit is 9.6. The facility also has three indoor electrical rooms. Two of these require cooling, provided by five CRAC units.

II. Electricity Use Characteristics

The facility's end-use of electricity is shown below in Table 1 and Chart 2. "Other" was calculated by subtracting all of the measured and calculated data from the "Whole Building" power. Eighty-two percent of the power goes to energizing the data center: 51 percent for server and related equipment loads, and 31 percent to the cooling equipment. The power used by the HVAC equipment is a significant amount of the whole building power consumption; it is where energy efficiency opportunities can be taken. A 15 to 50 percent reduction in HVAC electricity use can be achieved. This corresponds to 190 to 633 kW of electricity savings. Lighting power was calculated based on a design watt per square foot value. "Other" contributes 13 percent of the whole building load; it includes items such as losses of power in the electrical equipment, office plug loads, electrical room cooling equipment, and other office space conditioning equipment.

The power energizing the data center floor remains relatively constant throughout the day. This can be verified by examining Chart 4, which shows the power drawn by the server loads in a section of the data center over a period of five days. The server load power varied $\pm 1.6\%$ from the average of 428 kW. This indicates that power drawn remains the same regardless of the peak time usage of the servers and supporting equipment. The increase in power drawn midday of January 4 is due to the addition of more equipment by customers.

Table 1. End-Use of Electricity

Description	Electricity Consumption (kW)	% of Whole Building	Square Feet (sf)	Watts / sf
Whole Building	4,000	--	118,700	34
Data Center Server Load	2,040	51	60,400	34
Data Center CRAC Units	1,240	31		20
Office Space Conditioning	29	1	58,300	0.5
Lighting	173	4	118,700	1.5
Other	423	13	118,700	4.4

Chart 2. Facility 2 Electricity End-Use




Chart 3. Facility 2 Operating Energy Densities





Chart 4. Constant Power Drawn by Server Loads in Data Center




III. Electricity Use Diversity

In determining the size of the equipment needed in a data center facility, designers, in most cases, use an energy density value expressed in watts per square foot (W/sf). The type and number of server equipment in the data center are difficult to estimate, designers therefore oversize electrical equipment, so that a lack of capacity will not be a concern. Table 2 shows the operating conditions of the facility in comparison with the designed conditions. Also, an extrapolated value was calculated to determine what the operating W/sf would be if the data center were fully loaded.

Table 2. Server Load Diversity Factor

Measured W/sf	34
Extrapolated Full Load W/sf	52
Design W/sf	65
Diversity Factor (Measured / Design)	0.52
Diversity Factor (Extrapolated / Design)	0.80

Chart 5. Server Load Density



The data center was only 65 percent occupied at the time of monitoring. Thus 65 percent of the 60,400 sf area was used in calculating the extrapolated data center server load. The extrapolated diversity factor of 0.80 demonstrates that when the data center is fully occupied, it will operate at 80 percent of the expected or designed load.

**Appendix IV. “Data Center Energy Characterization Study,
Facility 3”; PG&E Case Study, Rumsey Engineers, Inc., 2001**

**DATA CENTER ENERGY
CHARACTERIZATION STUDY
SITE REPORT**

DATA CENTER FACILITY 3

FEBRUARY 2001

I. Review of Site Characteristics

Facility

Facility 3 is a two-story 42,000 square foot (sf) data center facility located in Silicon Valley, California. This facility provides co-location service, which is an unmanaged service that provides rack space and network connectivity via a high-capacity backbone. The building houses 25,000 sf of data center space, 4,580 sf of office/support space, and 900 sf of electrical room space. 59 percent of the building is data center space. During the monitoring period, 85 percent of the data center was occupied. Combined office and support space accounts for 11 percent of the building. The data center's environmental system operates 24 hours a day year-round. The users of the co-location space require full access to and control of their caged space 24 hours a day.




Customer Cage

Electrical Equipment and Backup Power System

The electricity use at Data Center Facility 3 was monitored from January 15, 2001 through January 22, 2001. Facility 3 has a PG&E service feed of 2,000 kVA. Of this, an average of 1,780 kVA is being used. The site drew an average of 1,760 kW over the monitoring period. Note that the power drops to zero on January 17, 2001 at 1:00 p.m. due to a rolling blackout. Data center floor power was not disrupted during the one and a half hours of the blackout. Disregarding the power outage, the load varied $\pm 4.5\%$ from the average, and the load factor was 0.97.

Chart 1. Facility 3 Whole Building Power Consumption



This site utilizes Liebert and PowerWare uninterruptible power supplies (UPS) to condition the utility power in order to shield critical loads on the data center floor from disturbances. In addition, a constant delivery voltage to the data center is maintained. Anytime the voltage drops below 480, the UPS systems' batteries feed in the necessary voltage to maintain 480 volts. The UPS converts AC current and stores it as DC current. When the voltage is needed, it is converted back to AC current. In the event of a power loss, 3 Onan diesel generators each provide 1,500 kW for backup. The generators with a full tank of diesel fuel can provide electricity to the building for approximately 8 hours. The data center is designed for N+1 redundancy on the electricity circuit level.



UPS System

Cooling System

The rooftop mechanical pad includes a farm of air-cooled condensers for the CRAC (computer room air conditioning) units, two package units for office space conditioning, and two exhaust fans.



Condenser Units

The data center is on a raised floor, through which cooling air is circulated via the CRAC units. Twenty-eight CRAC units serve the space: 22 Lieberts rated at 22 tons each and 8 Data Aire units rated at 10 tons each. The room temperature and relative humidity are maintained at an average of 65 °F and 40%, respectively.



Data Center CRAC Unit

Facility 3's electrical rooms are dispersed throughout the building. The DC power and transfer switch electrical rooms produce a high heat load. Thus they are cooled by two Liebert and four Data Aire CRAC units. These two rooms consist exclusively of equipment to power data center floor equipment. A temperature of 66 °F and relative humidity of 45% are maintained in the electrical rooms. The Liebert cooling units have a compressor EER (energy efficiency rating) of 11.7.

II. Electricity Use Characteristics

The facility's end-use of electricity is shown below in Table 1 and Chart 2. "Other" was calculated by subtracting all of the measured data from the "Whole Building" power. Eighty-eight percent of the power goes to energizing the data center: 59 percent for server and related equipment loads, 4 percent to the DC power equipment, and 25 percent to the cooling equipment. The DC power equipment is considered as part of the data center load since it is located on the data center floor. The HVAC electricity use is a significant amount of the whole building power consumption and is where energy efficiency improvements can be made. A 15 to 50 percent reduction in HVAC electricity use can be achieved. This corresponds to 66 to 221 kW of electricity savings. Lighting was calculated based on a design watt per square foot value. "Other" contributes 10 percent of the whole building consumption; it includes items such as losses of power in the electrical equipment, office plug loads, and office space conditioning.

Table 1. End-Use of Electricity

Description	Electricity Consumption (kW)	% of Whole Building	Square Feet (sf)	Watts / sf
Whole Building	1,760	--	42,000	42
Data Center Server Load	1,040	59	25,000	44
DC Power Equipment	68	4		
Data Center and Electrical Room Cooling	442	25	25,900	17
Lighting	42	2	42,000	1.0
Other	172	10	42,000	4.1



Chart 2. Facility 3 Electricity End-Use

Chart 3. Facility 3 Operating Energy Densities




III. Electricity Use Diversity

In determining the size of the equipment needed in a data center facility, designers, in most cases, use an energy density value expressed in watts per square foot (W/sf). The type and number of server equipment in the data center are difficult to estimate; designers therefore oversize electrical equipment, so that a lack of capacity will not be a concern. Table 2 shows the operating conditions of the facility in comparison with the designed conditions. Also, an extrapolated value was calculated to determine what the operating W/sf would be if the data center were fully loaded.

Table 2. Server Load Diversity Factor

Measured W/sf	44
Extrapolated Full Load W/sf	52
Design W/sf	75
Diversity Factor (Measured / Design)	0.59
Diversity Factor (Extrapolated / Design)	0.69

Chart 5. Server Load Density



The data center was 85 percent occupied at the time of monitoring. Thus 85 percent of the 25,000 sf area was used in calculating the extrapolated data center server load. The extrapolated diversity factor demonstrates that the data center is fully loaded, it will operate at 69 percent of the expected or designed load.

**Appendix V. “Data Center Energy Benchmarking Case Study,
Facility 4”; Federal Data Center Facility, Rumsey Engineers, Inc.;
LBNL - TengFang Xu, 2003**

DATA CENTER ENERGY BENCHMARKING CASE STUDY

FEBRUARY 2003



DATA CENTER FACILITY 4

Sponsored by



**Lawrence Berkeley National Laboratory
Environmental Energy Technologies Division**

Prepared By:



**99 Linden Street
Oakland, CA 94607
(510) 663-2070**

Acknowledgements

Rumsey Engineers and Lawrence Berkeley National Laboratory (LBNL) would thank the anonymous facility managers, engineers, and technicians for their generous assistance and cooperation in this data center benchmarking project. We would also thank Christine Condon of PG&E for providing monitoring equipment on short notice. This final report is edited by Tengfang Xu, the Project Manager of LBNL.

Disclaimer

Neither LBNL nor any of its employees makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any data, information, method, product or process disclosed in this document, or represents that its use will not infringe any privately-owned rights, including but not, limited to, patents, trademarks, or copyrights.

Contents

I.	EXECUTIVE SUMMARY.....	1
II.	DEFINITIONS	2
III.	SITE OVERVIEW	4
IV.	ENERGY USE.....	7
	UNINTERRUPTIBLE POWER SUPPLY	7
	DATA CENTER 4.1: DATA EQUIPMENT	8
	DATA CENTER 4.1: COOLING SYSTEM.....	9
	DATA CENTER 4.1: LIGHTING	15
	DATA CENTER 4.1: SUMMARY MEASUREMENTS AND METRICS	15
	DATA CENTER 4.5: DATA EQUIPMENT	19
	DATA CENTER 4.5: COOLING SYSTEM.....	20
	DATA CENTER 4.5: LIGHTING	21
	DATA CENTER 4.5: SUMMARY MEASUREMENTS AND METRICS	21
V.	ENERGY EFFICIENCY RECOMMENDATIONS	25
	GENERAL GUIDELINES FOR FUTURE CONSTRUCTION	25
	<i>Efficient Chilled Water System</i>	25
	<i>Air Management – HVAC Design</i>	27
	<i>Air Management – Rack Configuration</i>	27
	<i>Commissioning of New Systems and Optimized Control Strategies</i>	27
	<i>Lighting Controls</i>	28
	SPECIFIC RECOMMENDATIONS	28
	<i>Uninterruptible Power Supply</i>	28
	<i>Both Data Centers</i>	29

APPENDICES

- A: CHARTS OF MEASURED DATA
 - B: ESTIMATE OF DATA CENTER NON-CRITICAL LOAD
 - C: ESTIMATE OF DATA-AIRE CRAC UNIT PERFORMANCE AND EFFICIENCY
 - D: MARLEY TECHNICAL REPORT: THE APPLICATION OF COOLING TOWERS FOR FREE COOLING
-

I. Executive Summary

Rumsey Engineers and the Lawrence Berkeley National Laboratory (LBNL) have teamed up to conduct an energy study as part of LBNL's data center energy benchmarking efforts. This study is intended to provide measured information on energy and power use in data centers, and to help designers make better decisions about the design and construction of data centers in the near future. This report describes the outcomes of energy benchmarking in two data centers in Fresno, California, and the observations on potential opportunities in efficiency improvement. Measurements were conducted on-site from November 12 to 15, 2002, with the particular aim of determining the end-use of electricity power by infrastructures, computer equipment, and other components. The identity of data center owner and/or end-user is kept anonymous. The facility that houses both data centers is referred to as "Data Center Facility 4" throughout this report.

OBSERVATIONS AND RECOMMENDATIONS FOR ENERGY EFFICIENCY IMPROVEMENT

There are many opportunities for saving energy in this facility. The primary sources of inefficiency are the uninterruptible power supply (UPS), the Trane reciprocating chiller plant serving Data Center 4.1, and the computer room air-conditioning (CRAC) unit fans. This report concludes with ten recommendations that address these issues, four of which agree with those made in a May 2001 report presented to this facility through a previous Federal Energy Management Program (FEMP) study.¹

¹ "Assessment of Load and Energy Reduction Techniques (ALERT), Final Report" by the U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy, Federal Energy Management Program. Date of site visit, May 22-23, 2001.

II. Definitions

Data Center Floor Space	Gross footprint area of controlled data center space devoted to company/customer equipment. Includes aisles, caged space, cooling units, electrical panels, fire suppression equipment, and other support equipment. This gross floor space is what is typically used by facility engineers in calculating a computer load density (W/sf). ²
Data Center Server/Computer Load	Electrical power devoted to equipment on the Data Center Floor. Typically the power measured upstream of power distribution units or panels. Includes servers, switches, routers, storage equipment, monitors, and other equipment.
Critical Load	Electrical load of equipment that must keep running in the event of a power failure. Such loads are typically served by an Uninterruptible Power Supply (UPS), which uses a bank of batteries to support the load when the normal source of power fails. The batteries can support the load for only a short period. In some facilities the equipment is shut down gracefully and turned off until normal power returns. In other facilities a backup generator, typically diesel-powered, comes on-line and provides power for a longer period of time.
Data Center Cooling Power	Electrical power used to cooling equipment for the Data Center Floor Space.
Computer Equipment Occupancy	This is based on an estimate on how physically loaded the data centers are by computer and equipment footprints.
Measured Computer /Server Load Density	Ratio of actual measured Data Center Computer/Server Load in Watts (W) to the gross area (ft ² or sf) of Data Center Floor Space. Includes vacant space in floor area.
Projected Computer /Server Load Density	Ratio of projected Data Center Computer/Server Load in Watts (W) to square foot area (ft ² or sf) of the Data Center Floor Space if the Data Center Floor Space were fully

² Users look at watts per square foot in a different way. With an entire room full of communication and computer equipment, they are not so much concerned with the power density associated with a specific footprint or floor tile, but with larger areas and perhaps even the entire room. Facilities engineers typically take the actual UPS power output consumed by computer hardware and communication equipment in the room being studied (but not including air handlers, lights, etc.) and divide it by the gross floor space in the room. The gross space of a room will typically include a lot of areas not consuming UPS power such as access aisles, white areas where no computer equipment is installed yet, and space for site infrastructure equipment like Power Distribution Units (PDU) and air handlers. The resulting gross watts per square foot (watt/ft2-gross) or gross watts per square meter (watt/m2-gross) will be significantly lower than the watts per footprint measured by a hardware manufacturer in a laboratory setting.

	Floor Space if the Data Center Floor Space were fully occupied. The Projected Data Center Computer/Server Load Density is usually higher than actual measured density and can be calculated by multiplying the reciprocal of actual computer equipment occupancy.
Computer/Server Load Density per Rack Footprint	Measured Data Center Computer/Server Load in Watts (W) divided by the total floor area that the racks or equivalents occupy, or the rack “footprint”.
Cooling Load Density	The amount of cooling (tons) supplied to a given floor space (Ton/ft ² or sf)
Cooling Effectiveness Index	Ratio of electrical power devoted to cooling data center space to the electrical power used by computer and equipment. A lower number likely corresponds to more effective cooling.

III. Site Overview

On their web site (<http://datacenters.lbl.gov/What.html>), LBNL defines “data center” as follows:


“We define a data center as a special facility that performs one or more of the following functions:

- Store, manage, process, and exchange digital data and information;
- Provide application services or management for various data processing, such as web hosting Internet, intranet, telecommunication, and information technology.

We do not consider spaces that primarily house office computers including individual servers associated with work stations as data centers.”

Data Center Facility 4 is housed in a single building. In previous decades, according to staff, the data center area contained mostly large, mainframe style computers. This area represented about 7% of the total floor area of the facility. (See Figure 1.) The data center area was cooled independently of the rest of the building.

Figure 1. Facility Site Plan



Today, not all of the original data center area is used for data center purposes. Refer to Figure 3 for a simplified floor plan. Rooms 1 and 5, and the Network Room, now contain modern servers, storage drives, printers, and stand-alone PCs. Rooms 1 and 5 are sparsely occupied by current data center standards (Fig. 2), and contain a wider variety of equipment than is seen in a typical electronic commerce “server farm”, for example. There are relatively few rack-style servers.

Rooms 1 and 5 also have more than one use. Part of Room 1 is used for storage. Room 5 contains three desks that are regularly occupied by staff. These three desk stations aside, the measurement team observed relatively little foot traffic in Rooms 1 and 5.

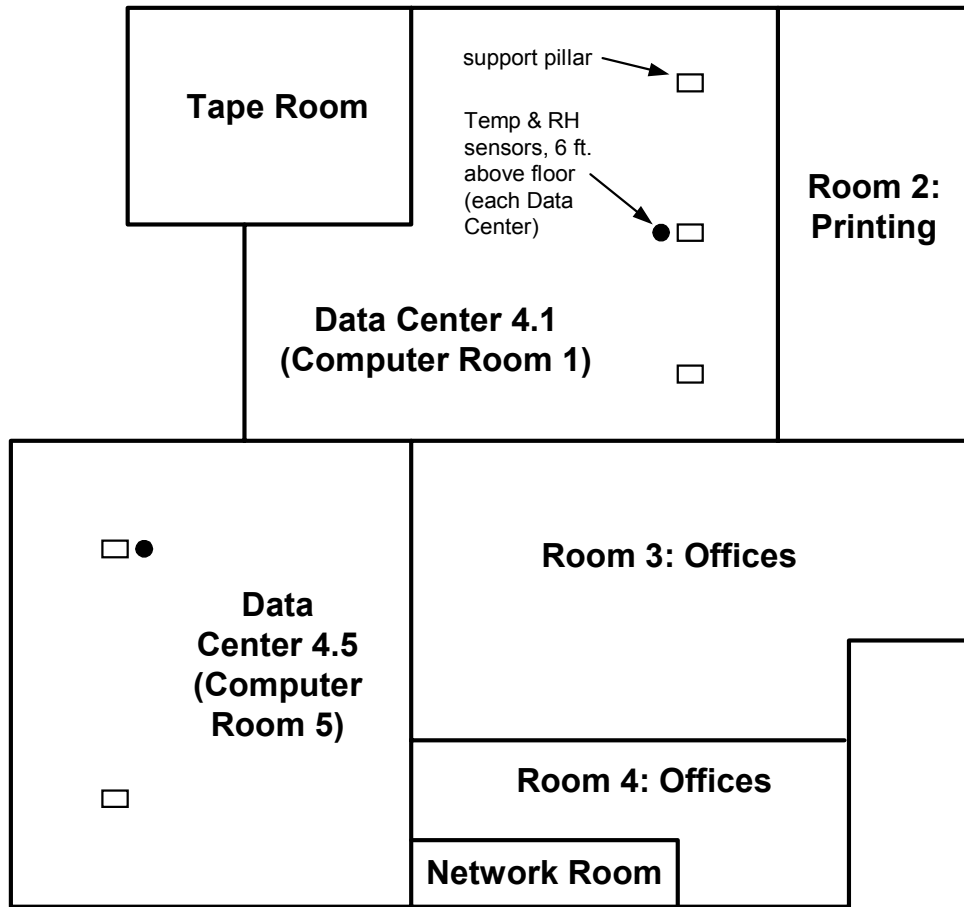
Figure 2. Example of Sparse Occupancy



Rooms 3 and 4 have been converted to office space. The Tape Room, Print Room (Room 2), and the Network Room continue to support data center operations, but are not included in this study. Rooms 3 and 4 now receive cooling from the same chilled water plant that serves the remainder of the building. The remaining rooms – 1, 2, 5, the Network Room and the Tape Room – continue to receive cooling from independent systems.

Rooms 1 and 5 are cooled in completely different manners, as described below. These two rooms are the focus of this report, and are hereafter referred to as Data Center 4.1 and Data Center 4.5. They are 8,900 square feet (sf), and 8,560 sf, respectively; this represents about 3% of the total building area.

Figure 3. Data Center Floor Plan



Data Centers 1 and 5 are cooled by dedicated Computer Room Air Conditioning (CRAC) units. The CRAC units are the only source of cooling for these two rooms. Outside air is provided indirectly by infiltration from adjacent rooms. The CRAC units in Data Center 4.1 use chilled water, normally provided by a dedicated plant comprised of four 100-ton reciprocating Trane water-cooled chillers. The CRAC units in Data Center 4.5 are split systems, with rooftop air-cooled condenser units.

Both data centers have raised floors, consisting of 2-foot by 2-foot tiles approximately 1 foot above the underlying slab. Most of the tiles are solid, but approximately 10% of them are perforated and distributed across the floor area. The CRAC units supply air to the underfloor space. Air rises through perforated tiles, and returns through grills located in the top of the CRAC units. All CRAC units are capable of humidity control.

A separate chilled water plant, consisting of two 750-ton Trane centrifugal chillers and a 1000-ton York centrifugal chiller, serves the rest of the building. This system was not monitored for this study.

IV. Energy Use

UNINTERRUPTIBLE POWER SUPPLY

A 450 kW Emerson Accupower model AP56 uninterruptible power supply (UPS) provides power to the critical loads in data center rooms 1 and 5, as well as the Print Room and Network Room. The UPS converts alternating current to direct current and charges a battery bank. Direct current from the batteries is converted back to alternating current and is fed to the data centers. In the event of a power outage the battery bank supplies power for about 30 minutes, which is long enough to permit a graceful shutdown of the critical computer loads. The UPS nameplate full-load efficiency is 92%.

The power supplied to and from the UPS was recorded for one day to determine the total critical computer load and the loss in the UPS itself.

Table 1. UPS Electrical Measurements

	Electrical Use ³	Units
UPS Input	127.0	kW
UPS Output	99.8	kW
UPS Loss	27.3	kW
UPS Efficiency	78.5	%

The measurements show that the UPS is operating well below its nominal capacity and its claimed efficiency. Refer to Appendix A for graphs of the recorded data.

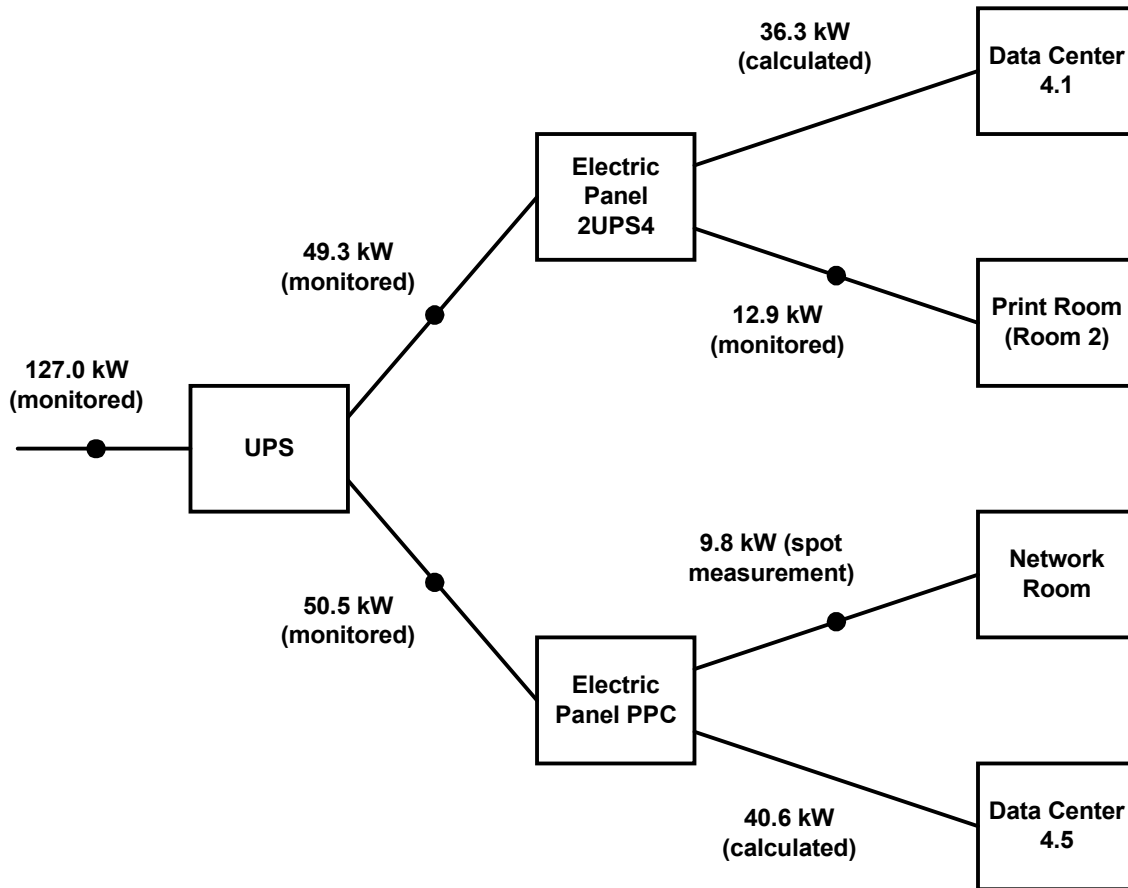
The electric panels serving the data centers are wired in such a way that it was easier to measure the total power of rooms 2 and 4, and obtain the power of data center rooms 1 (Data Center 4.1) and 5 (Data Center 4.5) by subtraction (Figure 5).

Figure 4. UPS



³ Average of 1-minute measurements taken during 11/14/02 - 11/15/02. Unless otherwise indicated, all readings of electric power in this report were made with a Summit Technology PowerSight PS3000 with 1000-amp clamp-on sensors.

Figure 5. Electrical Measurement Points



DATA CENTER 4.1: DATA EQUIPMENT

Data Center 4.1's computer equipment falls into two categories. All machines that must not immediately lose power in the event of a utility outage are collectively termed the "critical load". This equipment includes servers and their associated terminals (consisting of a monitor and keyboard), printers, and tape drives. This critical load is powered by the UPS. The total critical load was measured directly at two points immediately downstream of the UPS (Figure 5).

The remainder of the equipment (all of it PCs) is termed "non-critical". Rather than attempt to measure the non-critical load directly, an inventory of the non-critical equipment was made and the load was estimated. For details, see Appendix B. As shown in Table 2, the non-critical load is relatively small. In the summaries that follow, the total estimated load – 36.7 kW – is used in all calculations and the distinction between critical and non-critical is no longer made.

Table 2. Data Equipment Loads

	Electrical Use	Units
Critical	36.3 ⁴	kW
Non-Critical	0.4	kW

Figure 6. Servers in Data Center 4.1



DATA CENTER 4.1: COOLING SYSTEM

CHILLED WATER PLANT

Under normal operating conditions, the CRAC units in Data Center 4.1 receive chilled water from a chilled water plant that is dedicated to the data center area. The chiller plant contains four 100-ton reciprocating Trane water-cooled chillers. Only one of the four chillers (Chiller 3) was operating during the monitored period.

⁴ Average of 1-minute measurements taken 8:53 to 9:24 on 11/15/02.

The facility is equipped with a computer system for monitoring the function of both chilled water plants, but it is not capable of automatic control of the dedicated Trane chiller plant. Turning the chillers, pumps, valves and other devices on and off is done manually.

The plant has two 15-hp chilled water pumps. Only one of the two pumps (Pump 1) was operating during the period of measurement. The chilled water pump motors have continuously variable speed control, but the speed must be set manually – there is no automatic control.

There are two 15-hp condenser water pumps, of which one (Pump 3) ran during the monitored period. Like the chilled water pumps, the motors are equipped with manual speed control.

Chilled water pipes connect the two chiller plants, so the main plant can provide cooling to Data Center 4.1 if the data center plant goes off-line.

The data center chilled water plant has a pair of closed circuit, indirect cooling towers. Each cooling tower has two constant-speed fan motors. One motor drives two fans, and the other motor drives the remaining fan. Cooling tower 2 ran during the monitored period; cooling tower 1 did not. The condenser water is not mixed with the cooling tower water; it remains in closed pipes, while the cooling tower water wets the outside surface of the pipes. During the first stage of cooling, the fans remain off and a pump in the cooling tower sprays water over the pipes. For additional cooling both fan motors will turn on and off together as needed to meet the condenser water set point.

The use of closed-circuit cooling towers eliminates the need for the treatment chemicals normally used in condenser water. This allows the condenser water to be used as reheat water for the CRAC units (Figure 7). The reheat loop includes a reheat boiler, hot water coils in the CRAC units, and the condenser of Chiller 1. This arrangement is an energy-saving measure; it allows using waste heat from the chiller to warm the reheat water.

Typically the reheat boiler is not operated, though a reheat pump is normally operating.⁵

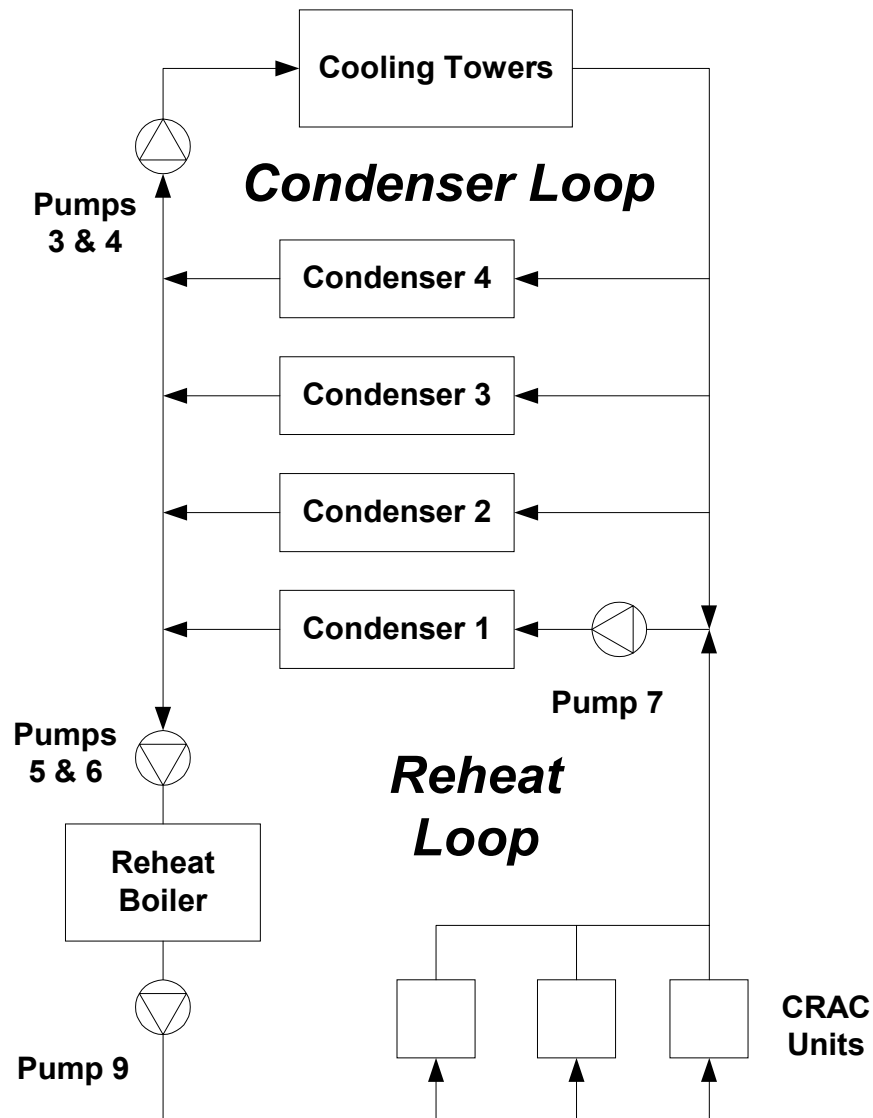
Condenser water is pumped through the condensers of all four chillers, regardless of the number of chillers operating.

The plant is plumbed to allow waterside economizing, called “Winter Cooling” in this facility. There is no heat exchanger; instead, the water in the chilled water loop is diverted directly to the cooling towers when waterside economizing is desired.⁶ This plant is unusual because the same water can be used for the condenser loop, the reheat loop, and the chilled water loop.

⁵ The reheat boiler hasn’t been on in the last 3 years, according to maintenance crew, yet a reheat pump (Pump 7) still runs continuously.

⁶ According to the maintenance crew, this feature hasn’t been used in recent years.

Figure 7. Simplified Schematic of Reheat and Condenser Loops in Data Center Chiller Plant



Note: This diagram shows only the condenser and reheat loops; it does not include the chilled water loop.

Figure 8. Measurement of Chilled Water Flow



Power consumption, flow, and chilled water temperatures⁷ were measured in the chiller plant over two days; see Table 3. The outside air temperature, relative humidity, and chiller measurements did not vary much during the monitored period. Appendix A contains graphs of the recorded data.

CRAC UNITS

The chilled water system serves water-cooled CRAC units in the data centers.


The CRAC units in Data Center 4.1 are constant-speed air handlers, equipped with electrically powered humidifying units. The CRAC units are fed three types of water – chilled water for cooling and dehumidifying, reheat water for reheating the air after dehumidifying; and deionized water for humidification. There are a total of six identical CRAC units in this data center, but only three were operating during the period of measurement. The units are turned on and off manually by the maintenance crew. According to facility staff, only three units are needed to satisfy the room set point during most of the year.

Air enters through filters in the top of the CRAC, and exits from the bottom to the underfloor space. The CRAC unit controls are set to maintain a return air temperature of $72^{\circ}\text{F} \pm 2^{\circ}\text{F}$, and a relative humidity of $50\% \pm 2\%$. These are relatively narrow ranges, and are a holdover from the time the data centers used mainframe computers. The actual

⁷ These were measured using a Summit Technology PowerSight PS3000 for electric loads, a Controlotron ultrasonic flow meter for chilled water flow, and a Pace Scientific Pocket Logger equipped with thermistors to measure the chilled water supply and return temperatures.

room air temperature and relative humidity were not measured directly by the investigation team. Temperature and relative humidity sensors are installed approximately in the middle of data centers (Figure 3). There is a data recorder for temperature and humidity, which was maintained by the facility engineers. The temperature and humidity remained fairly constant throughout the monitoring period. The temperature was around $72.5 \pm 0.5^\circ\text{F}$, while the relative humidity was around $49\% \pm 1\%$.

Figure 9. CRAC 7 in Data Center 4.1



Chilled water flow, and chilled water supply and return temperatures were measured at CRAC 7 for two days.⁸ A spot measurement was made of the electrical power usage of CRAC 7. The other two CRAC units in operation (9 and 10) were assumed to have the same cooling and electrical load. This is a reasonable assumption, as the room has a large volume and is sparsely occupied. Even though the computer equipment is not distributed evenly (there are areas of the room that are used for storage) it is unlikely that there are any “hot spots”. The ambient conditions for each CRAC unit are assumed to be essentially the same. The CRAC units are identical, constant-speed units, so the fan speed (and power) is assumed to be the same. The temperature and humidity set points were confirmed to be the same by examining the control panel on the front of each CRAC unit.

The spot measurements and the average of trended measurements are listed in Table 3. Refer to Appendix A for a graph of the CRAC 7 chilled water flow and temperature measurements over the entire monitoring period. The actual airflow through the CRAC unit was not measured.

⁸ Flow was measured with a Controlotron ultrasonic flow meter attached to the supply line, and temperatures were recorded with a Pocket Logger reading signals from thermistors tucked under the insulation of the chilled water supply and return lines.

Table 3. Cooling Equipment Electrical and Load Measurements

Equipment	Nominal Size	Spot / Monitored	Date(s)	Average Value
Chiller 3 Power	n/a	Monitored	11/13/02 - 11/15/02	46.6 kW
Chiller 3 Load	100 tons	Monitored	11/13/02 - 11/15/02	45.8 tons
Chilled Water Pump (Pump 1)	15 hp	Monitored	11/13/02 - 11/15/02	10.1 kW
Condenser Water Pump (Pump 3)	15 hp	Spot	11/13/02	13.2 kW
Reheat Water Pump (Pump 7)	5 hp	Estimated ⁹	--	2.9 kW
Cooling Tower 2, Fan 1	15 hp	Monitored	11/13/02 - 11/15/02	0.3 kW
Cooling Tower 2, Fan 2	7.5 hp	Spot	11/15/02	0.2 kW
Cooling Tower 2 Spray Pump	5 hp	Spot	11/13/02	4.0 kW
CRAC 7 Power	n/a	Spot	11/14/02	6.1 kW
Total CRACs Power (Units 7, 9, and 10)	--	Multiply by 3	--	18.3 kW
Amount of heat being removed from Data Center 4.1 by CRAC 7	n/a	Monitored	11/13/02 - 11/15/02	7.4 tons
Amount of heat being removed from Data Center 4.1 by CRAC Units 7, 9 and 10	--	Multiply by 3	--	22.1 tons

During the monitored period Chiller 3 delivered an average of 45.8 tons of cooling, approximately one-half of its nominal 100-ton capacity. This is only one-eighth of the total Trane chiller plant capacity.

The power consumption of the chilled water pump (Pump 1) is reasonable for a 15-hp motor, assuming 90% efficiency and a load factor of 81%. $[(15 \text{ hp})(0.746 \text{ kW/hp})(81\%)/(90\%) = 10.1 \text{ kW}]$ Assuming the condenser pump (Pump 3) has the same efficiency, it appears to be overloaded by about 6%, but this is not an uncommon motor service factor. $[(15 \text{ hp})(0.746 \text{ kW/hp})(106\%)/(90\%) = 13.2 \text{ kW}]$

The cooling tower fans cycled regularly during the monitored period. The fan power shown in Table 3 are average values. The total electric power consumption of the CRAC units in Data Center 4.1 is essentially the sum of the electric power draw of the fan motor

⁹ Gould E-Plus 5 hp motor; assumed 75% loaded, 95% efficient.

in each unit. This fan power represents 33% of the total HVAC power consumption. This power usage can probably be reduced; strategies are described in Section V—Energy Efficiency Recommendations, below.

DATA CENTER 4.1: LIGHTING

Lighting in Data Center 4.1 consists of T-8 fluorescent lamps with electronic ballasts. There are two lamps and one ballast per fixture. The room is equipped with occupancy sensors, though the facility operates 24 hours a day, 7 days a week. Foot traffic is frequent enough, and sensor delay long enough, that the lights rarely turn off, according to facility staff. The measurement team on this study supports this observation. All the lights were on during the monitored period. Counting the number of fixtures and multiplying by per-fixture wattage estimated the total lighting power:

$$(127 \text{ two-lamp fixtures}) \times (71 \text{ Watts/fixture}) = 9.0 \text{ kW}$$

The lighting power density is 1.0 Watts per square foot of gross floor area.

DATA CENTER 4.1: SUMMARY MEASUREMENTS AND METRICS


SUMMARY MEASUREMENTS

Table 4 brings together all the electrical measurements for the data center. The total UPS loss is split proportionally between the data centers, according to the critical load in each. Likewise, Data Center 4.1's share of the total electric power consumption of the dedicated chilled water plant is determined by the ratio of the measured cooling tons of the Data Center 4.1 CRAC units to the measured tonnage of the chiller plant.

Table 4. Summary of Electrical Measurements

Item	Remarks	Value	Percent
Computer Loads	Critical and non-critical	36.7 kW	33%
UPS Loss	Data Center 4.1 share of total UPS loss	10.0 kW	9%
HVAC – Air Movement	Fan power for CRAC Units 7, 9, 10	18.3 kW	16%
HVAC – Chilled Water Plant	Data Center 4.1 share of total chilled water plant power consumption	37.3 kW	34%
Lighting	--	9.0 kW	8%
Total Power Use	--	111.3 kW	100%

Figure 10. Energy Balance, Data Center 4.1



The total computer load is 33% of the data center electric power usage. The electricity used to cool the data center significantly outweighs this; pumping and cooling power accounts for 34% of the total, and air conditioning amounts to 16% of the data center electric power usage. UPS losses account for 9% of the data center electric power usage, and the lights account for the remaining 8%.

ELECTRICAL CONSUMPTION METRICS

Table 5 addresses the issue of data center load density. The most commonly used metric among mission critical facilities is the computer load density in Watts consumed per square foot (W/sf). However, even in a prototypical data center filled entirely with closely spaced racks of similar equipment, the choice of what to use as square footage is not always consistent between analysts, and can be a source of confusion.¹⁰ In the case of Data Center 4.1, there is wide variety of operating equipment in the room, for instances, servers in racks, stand-alone servers, stand-alone PCs, tape drives, printers.

¹⁰ See “Data Center Power Requirements: Measurements from Silicon Valley”, by Mitchell-Jackson, Koomey, Nordman, & Blazek, December 2001. It is available on the web at http://enduse.lbl.gov/Info/Data_Center_Journal_Article2.pdf.)

Less than 10% of gross floor areas are occupied by the equipment, which includes all types of computer equipment.

Table 5. Electrical Consumption Metrics

Metric	Value	Units
Data Center Gross Area	8,900	sf
Computer Equipment Area	800	sf
Computer Load Density based on Gross Area	4.1	W/sf
Current Computer/Equipment Occupancy	9	%

“Data Center Gross Area” is the entire floor area of the room, including storage areas and the main aisles. Per the Uptime Institute Definitions, this gross floor space is what facility engineers typically use when calculating a computer load density (W/sf).¹¹ “Computer Equipment Area” is the portion of the room occupied by all computer equipment, not including the narrow aisle space immediately around the equipment.

The total computer equipment load is 36.7 kW. The computer load density based on the gross area is 4.1 W/sf. Compared to typical 30 to 50 W/sf as seen in other facilities with relatively tight packing, this computer load density is rather low largely because of its sparse distribution. Typical data centers contain tall racks set close together in rows, with a narrow aisle between each row. Each slot in a rack can contain a board with one or more processors on it. For the purpose of estimating the maximum potential build-out of the data center, it is assumed that more of the same general type of equipment will be added to the footprints of data center floor space. This would theoretically result in the computer load density of no more than 45 W/sf, while the total computer load at 100% occupancy would increase up to 400 kW.

HVAC EFFICIENCY METRICS

Since the packing of data centers and computer types are site specific, a more useful metric for evaluating how efficiently the data center is cooled can be represented as a ratio of cooling power to computer power. Another metric is the “theoretical cooling load”. It is the sum of the computer, lighting, and CRAC fan motor loads. All of these loads equate to heat that must be removed from the room. Though there is a small amount of human activity, this is insignificant compared to the other loads.

¹¹ Users look at watts per square foot in a different way. With an entire room full of communication and computer equipment, they are not so much concerned with the power density associated with a specific footprint or floor tile, but with larger areas and perhaps even the entire room. Facilities engineers typically take the actual UPS power output consumed by computer hardware and communication equipment in the room being studied (but not including air handlers, lights, etc.) and divide it by the gross floor space in the room. The gross space of a room will typically include a lot of areas not consuming UPS power such as access aisles, white areas where no computer equipment is installed yet, and space for site infrastructure equipment like Power Distribution Units (PDU) and air handlers. The resulting gross watts per square foot (watt/ft²-gross) or gross watts per square meter (watt/m²-gross) will be significantly lower than the watts per footprint measured by a hardware manufacturer in a laboratory setting.

Chiller efficiency is usually presented as the ratio of chiller power at full load to the tons of cooling provided at full load, in units of kW/ton. HVAC system efficiency is similar, but it takes in to account the power consumption of all the HVAC system components – chiller, cooling tower, pumps, and air handlers.

Table 6. HVAC Efficiency Metrics

Metric	Average Value	Units
Cooling kW / Computer Load kW	1.5	--
Theoretical Cooling Load	18.2	tons
Cooling Provided by HVAC System	22.1	tons
Chiller Efficiency	1.0	kW/ton
HVAC System Efficiency	2.5	kW/ton

The “Cooling Effectiveness Index ” is 1.5 kW/kW (a lower number corresponds to more effective cooling). This indicates that in Data Center 4.1, for every unit of electric power used by computer equipment, 1.5 units of electric power is being used by cooling to maintain the conditions of the data center. This figure does not include the efficiency loss in the UPS, or the power required to keep the UPS cool.

In the case of Data Center 4.1 the theoretical load is 20% less than the measured cooling delivered by the HVAC system. There are several possible explanations for this. The non-critical load may be underestimated. Some heat may be entering the data center from adjoining rooms. Or, CRAC units 8 and 9 may be providing a little less cooling than CRAC unit 7. The chiller efficiency is 1.0 kW/ton. Though this is typical for an older reciprocating unit, newer water-cooled chillers are much more efficient. For the portion of the chilled water system that serves Data Center 4.1, plus the CRAC units in that room, the overall HVAC system efficiency is 2.5 kW/ton. An unusually large portion of the HVAC power consumption is due to the CRAC units, even though these units consist mainly of a fan and cooling coil, and do not perform a refrigeration process.

Strategies for improving efficiency are addressed later in this report.

DATA CENTER 4.5: DATA EQUIPMENT

As in Data Center 4.1, there are critical and non-critical computer loads. The critical load was measured directly, and the non-critical load was estimated. For details of the estimation, see Appendix B. As in Data Center 4.1, the non-critical load is relatively small. The total estimated load – 41.4 kW – is used in summary calculations and the distinction between critical and non-critical is no longer made.

Table 7. Data Equipment Loads

	Electrical Use	Units
Critical	40.6	kW
Non-Critical	0.8	kW

Figure 11. Servers in Data Center 4.5



DATA CENTER 4.5: COOLING SYSTEM

Data Center 4.5 contains ten identical split-system, air-cooled CRAC units. They are Data-Aire units, model DAAD-2034. See Appendix C for specifications. The compressor, evaporator, air delivery fan, and humidifier for each system are in a single housing (the “evaporator unit”) that stands inside Data Center 4.5. The units have two stages of cooling. They are capable of dehumidification; reheat is provided by waste heat from the compressor. The warm refrigerant is piped to the rooftop condenser units where it passes through a refrigerant-to-air heat exchanger, before returning to the evaporator unit. Four multi-speed fans in the condenser unit stage on as needed, to cool the refrigerant to the desired temperature.

Figure 12. Air-Cooled Condensers on Roof



During the period of measurement, only four of the ten evaporator units were running – CRAC units 23, 24, 25 and 28. The outside dry bulb air temperature varied between 50 °F and 60 °F during this time, and the outside air relative humidity varied between 80% and 100%. A chart of outside air conditions is included in Appendix A. The remaining evaporator units were shut off. Facility staff reports they are not needed to keep the room at the desired temperature. The CRAC units control to the same narrow standard as in Data Center 4.1: a return air temperature of 72 °F \pm 2 °F, and a relative humidity of 50% \pm 2%. The actual room air temperature and relative humidity were not measured directly by the investigation team. There is a data recorder for temperature and humidity within the data center space, which was maintained by the facility engineers. The temperature and humidity remained fairly constant throughout the monitoring period. The temperature was around 70 \pm 1°F, while the relative humidity was around 52% \pm 1%.

The display panel on CRAC 23 indicated it was in dehumidifying mode, unlike the other three operating CRAC units. The power consumption of the evaporator and condenser units of CRAC 23 was monitored for about 18 hours. A spot measurement was taken of the evaporator power for CRAC 24. The spot measurements and average of trended measurements are listed in Table 8. Refer to Appendix A for graphs of the measurements over the entire monitored period. The actual airflow through the CRAC units was not measured.

Table 8. Cooling Equipment Electrical Measurements

Equipment	Spot / Monitored	Date(s)	Average Value
CRAC 23, Evaporator Unit	Monitored	11/14/02 – 11/15/02	12.2 kW
CRAC 23, Condenser Unit	Monitored	11/14/02 – 11/15/02	0.7 kW
CRAC 24, Evaporator Unit	Spot	11/13/02	4.3 kW
Total Evaporator Unit Power, CRAC Units 23, 24, 25, 28	Sum	--	33.0 kW
Total CRAC Unit Power, Units 23, 24, 25, 28	Sum	--	34.4 kW

CRAC 23 consumed almost 3 times as much power as CRAC 24, and the consumption was steady during the monitored period. Based on the power consumption measurements, it appears that only the fan in CRAC 24 was operating, and the unit was not doing any active cooling. This assertion is supported by the analysis in Appendix C. If true, then the CRAC 23 compressor power is assumed to be $12.2 \text{ kW} - 4.3 \text{ kW} = 7.9 \text{ kW}$. (Evaporator Unit Power minus Fan Power equals Compressor Power.) Based on the analysis, it is assumed that two of the four CRAC units were operating like CRAC 23 (compressor operating), and the remaining two were operating like CRAC 24 (compressor not operating).

DATA CENTER 4.5: LIGHTING

The type of light fixtures and the number of operating hours in Data Center 4.5 are identical to Data Center 4.1, with a different number of fixtures. All the lights were on during the monitored period. The total lighting power is calculated by multiplying the number of fixtures by 71 Watts/fixture.

$$(120 \text{ two-lamp fixtures}) \times (71 \text{ Watts/fixture}) = 8.5 \text{ kW}$$

The lighting density is 0.99 Watts per square foot of gross floor area, very similar to the 1.02 W/sf value for Data Center 4.1.

DATA CENTER 4.5: SUMMARY MEASUREMENTS AND METRICS


SUMMARY MEASUREMENTS

Table 9 brings together all the electrical measurements for the data center. The total UPS loss is split proportionally between the data centers, according to the critical load in each.

Table 9. Summary of Electrical Measurements

Item	Remarks	Average Value	Percent
Computer Load	Critical and non-critical	42.6 kW	43%
UPS Loss	Data Center 4.5 share of total UPS	11.1 kW	12%
HVAC	Total power consumption of four CRAC units: 23, 24, 25, 28	34.4 kW	36%
Lighting	--	8.5 kW	9%
Total Power Use	--	96.7 kW	100%

Figure 13. Energy Balance, Data Center 4.5



The critical computer load served by the UPS, plus the non-critical computer load, amounts to 43% of the data center electric power usage. Cooling Power, including air movement, is the second largest consumer at 36% of the data center electric power usage. UPS loss accounts for 12% of the data center electric power usage, and the lights account for the remaining 9%.

ELECTRICAL CONSUMPTION METRICS

Unlike Data Center 4.1, computer equipment is spread uniformly over the entire floor area of Data Center 4.5. However, the equipment floor density is also very low – they occupy 600 ft² out of the total data center floor area of 8,560 ft². The aisles between equipment are wide, and the four work cubicles that have been set up do not interrupt the uniformity of computer placement. The cubicles were staffed continuously during the period of measurement, but the heat load represented by the occupants is negligible in comparison to the equipment. Given this arrangement, only the computer load density based on gross data center area is shown in Table 10.

Table 10. Electrical Consumption Metrics

Metric	Value	Units
Data Center Gross Area	8,560	sf
Computer Equipment Area	600	sf
Computer Load Density based on Gross Area	4.8	W/sf
Current Computer/Equipment Occupancy	7	%

The total computer load is 41.4 kW. The actual computer load density based on the gross area is 4.8 W/sf, which is well below current typical load densities largely due to rather sparse computer equipment occupancy. Based on information provided by the facility engineers, it is estimated that Data Center 4.5 is about 7% occupied. Assuming that more computer equipment of the same type would be added, and that the room would be built out to 100% occupancy in terms of footprints, the computer load density would theoretically be no more than 68 W/sf, while the total computer load would theoretically increase up to 590 kW.

HVAC EFFICIENCY METRICS

The theoretical cooling load is calculated in the same manner as for Data Center 4.1. Time constraints did not allow direct measurement of the cooling provided by the CRAC units. The total cooling provided was estimated after analyzing the specifications for the CRAC unit. See Appendix C for details.

Table 11. HVAC Efficiency Metrics

Metric	Average Value	Units
Cooling kW / Computer Load kW	0.8	--
Theoretical Cooling Load	19.1	tons
Estimated Cooling Provided by CRAC Units	16.8	tons
Estimated CRAC Unit Efficiency	1.4	kW/ton

The “Cooling Effectiveness Index ” is 0.8 kW/kW (a lower number corresponds to more effective cooling) in Data Center 4.5. This indicates that the cooling in this data center is

almost twice as effective as the cooling in Data Center 4.1 with the “Cooling Effectiveness Index ” of 1.5 kW/kW. Although water-cooled chiller systems are typically more efficient than air-cooled, the discrepancy in this case is due to a combination of the following likely reasons: 1) that the chilled water plant serving Data Center 4.1 is very much under loaded, resulting in inefficient operation; 2) that Data Center 4.1’s chilled water plant has old reciprocating chillers, which are less efficient than centrifugal, screw, or scroll chillers; 3) that computer equipment are much more evenly spread out in Data Center 4.5, which likely mitigates the cooling demand across the whole data center; and 4) that CRACs in Data Center 4.1 might have been overly used.

The theoretical cooling load in Data Center 4.5 is 19.1 tons, similar to Data Center 4.1’s theoretical cooling load of 18.2 tons. This is expected, given that the two data centers have an essentially similar scale of floor area, and contain similar types and quantities of computer equipment, lights, and CRAC unit fan motors. The estimated cooling provided by the CRAC units in Data Center 4.5 is 12% lower than the theoretical load the CRAC units are serving. As both numbers are estimates, this is reasonable agreement. The amount of cooling provided by the CRAC units, and their efficiency, were estimated by the method described in Appendix C.


V. Energy Efficiency Recommendations

GENERAL GUIDELINES FOR FUTURE CONSTRUCTION

Efficient Chilled Water System

Water-cooled chillers offer enormous energy savings over air-cooled chillers, particularly in dry climates, because they take advantage of evaporative cooling. Since the chiller is being cooled by lower temperature media, it can reject heat more easily, and does not have to work as hard. Though the addition of a cooling tower adds maintenance costs associated with the water treatment, we have found that the energy savings outweigh the maintenance costs. Within the options of water cooled chillers, variable speed centrifugal are the most energy efficient, because they can operate very efficiently at low loads. The graph below compares the energy performance of various chiller types.

Comparison of Typical Chiller Efficiencies over Load Range

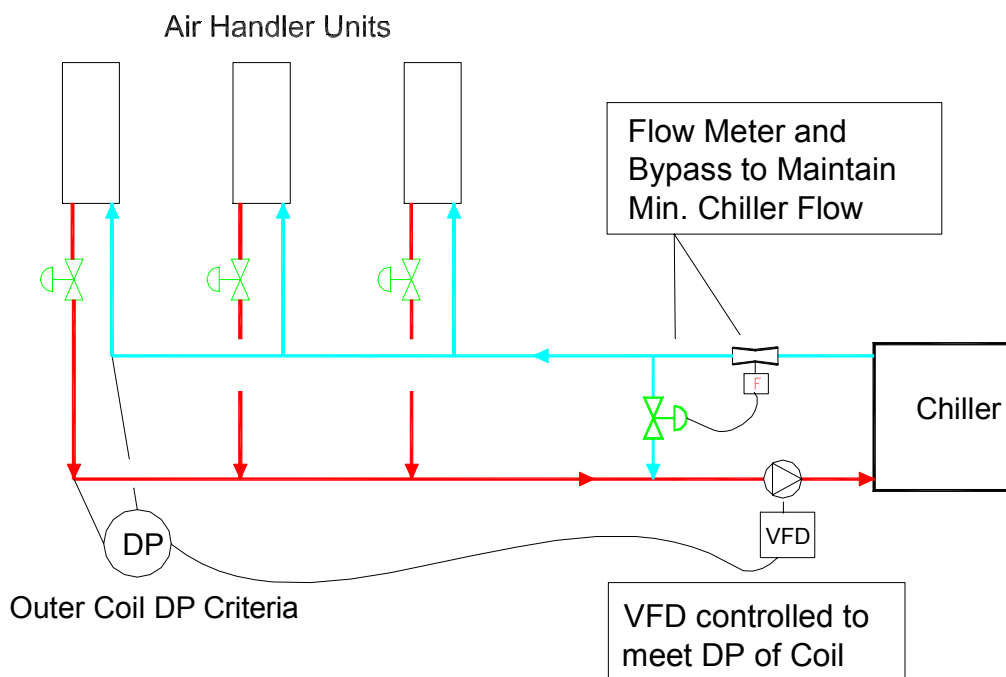


Chiller 1	250-Ton, Screw, Standard Efficiency, Air Cooled
Chiller 2	216 Ton, Screw, Water Cooled
Chiller 3	227-Ton, Centrifugal, Constant Speed, Water Cooled
Chiller 4	227-Ton, Centrifugal, Variable Speed, Water Cooled

Though there are efficient air cooled chillers, the larger size of water cooled chillers has resulted in more care given to efficiency and life cycle costs compared to air cooled chillers.

The selection of the auxiliary equipment, including the cooling tower, pumps, and pumping strategy should also be considered carefully. For example, variable speed fans on cooling towers allow for optimized cooling tower control. Premium efficiency motors and high efficiency pumps are recommended, and variable speed pumping is a ripe opportunity for pump savings. Variable pumping strategies can be achieved in a primary/secondary scheme, where the primary pumps operate at constant speed and directly feed water to the chiller, and the secondary pumps are variable speed and serve the air-handling units. A more energy efficient scheme is primary-only variable speed pumping strategy. Pumping savings are based on the cube law: pump power is reduced by the cube of the reduction in pump speed, which is directly proportional to the amount of fluid pumped.

A primary only variable pumping strategy must include a bypass valve that ensures minimum flow to the chiller, and the use of two-way valves at the air-handling units in order to achieve lower pumping speeds. The control speed of the bypass valve should also meet the chiller manufacturers recommendations of allowable turndown, such that optimum chiller efficiency is achieved.¹² The diagram below describes the primary-only variable speed pumping strategy.



¹² This basically means that the flow through the chiller should be varied slow enough such that the chiller is able to reach a quasi-steady state condition and able to perform to its maximum efficiency.

Air Management – HVAC Design

The standard practice of cooling data centers employs an underfloor system fed by CRAC units. There are a number of potential problems with such systems: an underfloor system works on the basis of thermal stratification. This means that as the cool air is fed from the underfloor, it absorbs energy from the space, warming up as a result, and rises. In order to take advantage of thermal stratification, the return air must be collected at the ceiling level. CRAC units often have low return air grills, and are therefore, simply recirculating cool or moderately warmed air. Furthermore, they are often located along the perimeter of the building, and not dispersed throughout the floor area, where they can more effectively treat warm air. One alternative is to install transfer grills from the ceiling to the return grill. Another common problem with underfloor supply is that the underfloor becomes congested with cabling, increasing the resistance to air flow. This results in an increase in fan power use. A generous underfloor depth is essential for effective air distribution (we have seen 3 feet in one facility).

An alternative to underfloor air distribution is high velocity overhead supply, combined with ceiling height return. A central air handling system can be a very efficient air distribution unit. Design considerations include using VFDs on the fans, low-pressure drop filters, and coils. An additional advantage of a central air handling system is that it can be specified with an economizer function. With the favorable climate in the Bay Area, economizing can reduce the cooling load for a majority of the hours of the year.

Another common problem identified with CRAC units is that they are often fighting each other in order to maintain a constant humidity set point. Not only is a constant humidity set point unnecessary for preventing static electricity (the lower limit is more important), but also it uses extra power. A central air-handling unit has a better ability to control overall humidity than distributed CRAC units.

Air Management – Rack Configuration

Another factor that influences cooling in data centers is the server rack configuration. It is more logical for the aisles to be arranged such that servers' backs are facing each other, and servers' fronts are facing each other. This way, cool air is drawn in through the front, and hot air blown out the back. The Uptime Institute has published documents describing this method for air management.¹³ Our observations of both data centers showed an inconsistent rack configuration.

Commissioning of New Systems and Optimized Control Strategies

Many times the predicted energy savings of new and retrofit projects are not fully realized. Often, this is due to poor and/or incomplete implementation of the energy efficiency recommendations. Commissioning is the process of ensuring that the building systems perform as they were intended to by the design. Effective commissioning

¹³ <http://www.upsite.com/TUIpages/whitepapers/tuiaaisles.html>

actually begins at the design stage, such that the design strategy is critically reviewed. Either the design engineer can serve as the commissioning agent, or a third party commissioning agent can be hired. Commissioning differentiates from standard start-up testing in that it ensures systems function well relative to each other. In other words, it employs a systems approach.

Many of the problems identified in building systems are often associated with controls. A good controls scheme begins at the design well. In our experience, an effective controls design includes 1) a detailed points list, with accuracy levels, and sensor types, and 2) a detailed sequence of operations. Both of these components are essential for successfully implementing the recommended high efficiency chilled water system described above. Though commissioning is relatively new to the industry, various organizations have developed standards and guidelines. Such guidelines are available through organizations like the Portland Energy Conservation Inc., at www.peci.org, or ASHRAE, Guideline 1-1996.

Lighting Controls

Lighting controls such as occupancy sensors may be appropriate for areas that are infrequently or irregularly occupied. If 24-hour lighting is desired for security reasons, scarce lighting can be provided at all hours, with additional lighting for occupied periods.

SPECIFIC RECOMMENDATIONS

Uninterruptible Power Supply

UPS.1. Consider Replacing the UPS

The UPS nameplate indicates an efficiency of 92%; the actual efficiency was measured at 78.5%. This is likely a result of unloading the unit to the point where is operating at only 28% of its nominal capacity of 450 kW. As a result, 27.3 kW of electricity is continuously converted into waste heat. In addition, this heat has to be removed from the UPS room by a dedicated package-cooling unit. If we assume the efficiency of this cooling unit is constant at 1.0 kW/ton (a reasonable number for an efficient package unit), and electricity cost is \$0.10/kWh, then the annual cost of this UPS inefficiency is about \$30,000. A new, appropriately sized UPS can run at 93% efficiency. Given the same preceding assumptions, such a new unit would reduce the annual inefficiency cost by two-thirds.

UPS.2. Consider Placing More Load on the UPS

Another approach is to place more of the facility's computer equipment on the UPS, to bring its load closer to its nominal 450 kW. Assuming the UPS is still capable of operating at or near its nameplate efficiency of 92%, this strategy would confer the

advantages described in measure UPS.1, as well as protecting more equipment in the event of a power outage.

Both Data Centers

These recommendations apply to both data centers. Recommendations specific to each data center follow.

DC 4.0 - 1. Run Fewer CRAC Units

The three operating CRAC units in Data Center 4.1 are lightly loaded. The cooling delivered by CRAC 7 was measured at 7.4 tons, and CRAC units 9 and 10 are assumed to be similarly loaded. The fans in each of these units run at constant speed and use a total of 18.3 kW, which is a relatively high 33% of total HVAC power consumption. If Data Center 4.1 can maintain temperature uniformity by running only two CRAC units, each at 11 tons, then the cost of running the third fan can be saved. If electricity costs \$0.10 per kWh, then this will avoid about \$5,350 a year in fan power.

Data Center 4.5 is operating four CRAC units. It appears that two of them are lightly loaded (8.4 tons), and that two of them are running their fans only; i.e., the compressors do not appear to be on). The fans are constant speed and draw a continuous 4.3 kW. If electricity costs \$0.10 per kWh, then each fan costs about \$3,750 per year to operate. If the temperature set point can be maintained with fewer operating CRAC units, each working proportionally harder to provide the same overall amount of cooling, then this fan power can be saved.

Important Note: Care should be taken to account for the fire blocking underneath the floor. The raised floor areas are physically divided in to zones, which limits the options when turning CRAC units off. Each zone that contains computer equipment will likely need at least one operating CRAC unit.

DC 4.0 – 2. Revisit Temperature and Relative Humidity Setpoints

The temperature and relative humidity standards in the data centers – $72^{\circ}\text{F} \pm 2^{\circ}\text{F}$, and $50\% \pm 2\%$ – are relatively narrow, and are a holdover from the time the data centers used mainframe computers. Maintaining these standards requires greater energy use. The equipment currently in the data centers can likely operate without any problems with more relaxed temperature and humidity standards.

DC 4.0 – 3. Rearrange Floor Tiles for More Effective Cooling

Some of the tiles in the data centers are perforated, to allow the cooling air to rise from the space under the floor. The investigation team noticed that many of the perforated tiles could be rearranged to more efficiently cool the computer equipment. In particular, perforated tiles should be placed in front of server racks, not behind. The cooling fans inside the servers typically draw air from the front of the rack and eject it out the back.

Directing the cooling air that is coming from the floor to rise in front of the rack will provide the optimum cooling effect.

DC 4.0 - 4. Reduce Lighting

Data Center 4.1 is large and has only intermittent visitors, yet it uses some 9 kW of lighting energy almost constantly. Consider any of the following measures:

- a) Reduce the lights-on interval of the existing occupancy sensors.
- b) Add more, narrow-view occupancy sensors, and wire them to specific lighting zones, so that only needed lights come on as people move through the room.
- c) Add task lighting in appropriate areas, and disable a portion of the overhead lights.

The same strategy for reducing lighting energy applies to Data Center 4.5. In addition, accommodation for the regular occupants is needed. Therefore, consider introducing task-specific, efficient lighting, such as electronic-ballast compact fluorescent desk lamps. Newer models have very high quality, flicker free light, and draw as little as 18 Watts.

Data Center 4.1

DC 4.1 - 1. Turn Off the Reheat Pump (Pump 7)

The CRAC units in Data Center 4.1 are able to dehumidify the air in the room by overcooling it. To bring the air temperature back to the desired level, reheat water is required. During the two days of measurement the outside air relative humidity was continuously over 80%, and at times reached 100%, yet the reheat system did not turn on. According to maintenance staff, the reheat boiler hasn't run in the last 3 years. These observations indicate that excess humidity is not a problem in the data center.

The 5-horsepower reheat pump (Pump 7) runs continuously, and needlessly. If we assume the motor is 75% loaded, and 95% efficient, then it draws a constant 2.9 kW. If electricity costs \$0.10 per kWh, this pump costs about \$2,500 a year to run. It should be turned off.

DC 4.1 - 2. Use Chilled Water From the Main Plant

The four 100-ton reciprocating Trane chillers are older and almost certainly less efficient than the large centrifugal chillers in the main plant. Chilled water pipes already run from the main plant to the smaller plant, and the CRAC units can be run directly with chilled water from the main plant. We recommend that this strategy be used whenever the main plant is running. The small Trane chiller plant should be shut off during these periods. We did not measure the efficiency and operating conditions of the main plant so we are unable to calculate cost savings, but they should be significant. For the periods where the small Trane plant must be run, consider the following measures:

DC 4.1 - 3. Use Free (“Winter”) Cooling

The small Trane chiller plant is designed to allow waterside economizing, also known as “free cooling”. Signs in the chiller room indicate that the term “winter cooling” is also used. This strategy uses the cooling towers to produce chilled water directly when the outside air conditions are favorable; i.e., low humidity. The chillers are shut off during these periods. If normally only one chiller operates, and the chiller normally draws 50 kW, and the cost of electricity is \$0.10 per kWh, then this strategy saves \$5/hour of chiller energy.

The cooling towers in this plant (as well as all the other components) are sized for a much larger load than currently exists in the data centers served by the plant. The cooling towers can likely make sufficient chilled water for the data centers for a large part of the year, given Fresno’s dry climate. The maintenance crew indicated that free cooling hasn’t been used in a long time. Consider re-implementing this scheme. See Appendix D for a technical paper that describes the strategy in detail.

DC 4.1 - 4. Add a Variable Speed Drive to the Cooling Tower Fans

As shown by the chart in Appendix A, Fan 1 in Cooling Tower 2 is cycling on and off frequently (approximately once per hour). The fan motor uses 2 kW when it is on, but the average consumption over the monitored period was only 0.3 kW. In other words, the motor has a duty cycle of about 15%. A variable speed drive will allow the fan to run continuously and more slowly, extending the life of the equipment. If the cooling towers are called upon to handle a larger load in the future, the variable speed drive will offer significant power savings as well.

DC 4.1 - 5. Feed Condenser Water Only Through Running Chiller

The piping and valve arrangement in the small Trane chiller plant allows the condenser water to be sent selectively through only the chiller that is operating. The condenser water was observed to be flowing through all four chillers simultaneously. Shutting the valves to the 3 non-running chillers will provide several benefits:

- a) The condenser pump is constant speed. Confining the condenser water flow to one condenser instead of four will increase the resistance to flow. The pump will “ride its curve”, and deliver less flow at a higher pressure. Counter-intuitively, this will actually reduce its power consumption.
- b) Even though the total gallons per minute (gpm) will be lower, the gpm through the single operating chiller will be higher than it was before, because the entire flow is going through a single condenser instead of four. The temperature of the condenser water leaving the chiller will be lower than it was before, which will help the chiller operate more efficiently.
- c) If the chiller is rejecting the same amount of heat as before, but the total condenser water flow has been reduced, then the temperature of the condenser water going to

the cooling tower will be higher than before. This will allow the cooling tower to more easily reject heat to the outdoors, which in turn allows the cooling tower fans to run less often, saving energy.

DC 4.1 - 6. Feed Chilled Water Only Through Running Chiller

The four Trane chillers are connected in series-parallel. The evaporators for Chiller 1 and 2 are piped in series, as are the evaporators for Chillers 3 and 4. Each of these pairs is then piped in parallel. As in the case of the condenser water, the piping and valve arrangement allows the chilled water to be sent only through the chiller that is running. We observed that the chilled water is flowing through both Chiller 3 and 4.

Adjusting the valves to route the chilled water only through the running chiller will reduce the resistance to flow, lowering the load on the chilled water pump. Since the chilled water pump has a variable speed control, the pump speed can be reduced accordingly, saving energy.

Data Center 4.5

DC 4.5 - 1. Investigate Dehumidifying Mode of CRAC 23

Observations in Data Center 4.1 indicate that excess humidity is not an issue in either of the data centers, yet CRAC 23 in Data Center 4.5 displayed “dehumidifying” on its front panel. It is not uncommon for CRAC unit controls to fail in dehumidifying mode. This causes unnecessary and intensive power use, and can go undetected for a long time. Measurements of CRAC 23 power did not reveal excessive usage; however, this status message bears investigation.


Several of the above recommendations – DC 4.0-4 (reduce lighting), DC 4.1-1 (turn off reheat pump), and DC 4.1.2 (use chilled water from main plant) – agree with those made in a May 2001 report presented to this facility through a Federal Energy Management Program (FEMP) study.¹⁴

¹⁴ “Assessment of Load and Energy Reduction Techniques (ALERT), Final Report” by the U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy, Federal Energy Management Program. Date of site visit, May 22-23, 2001.


Appendix A

Charts of Measured Data


Data Center Facility 4 Outside Air Conditions




Data Center Facility 4 UPS Input Power




Data Center Facility 4
Panel 2UPS4 Power
(Computer Rooms 1 & 2, served from UPS)




Data Center Facility 4
Panel PPC Power
(Computer Rooms 4 & 5, served from UPS)




Data Center Facility 4
Computer Room 2 Electrical Load (served from UPS)




Data Center Facility 4 Data Center Loads



Data Center Facility 4
Trane Chiller 3 Electrical Load




Data Center Facility 4 Trane Chiller 3 Operating Conditions




Data Center Facility 4


Chilled Water Pump 1 Electrical Load




Data Center Facility 4
Cooling Tower 2, Fan Motor 1 Electrical Load



Data Center Facility 4
CRAC 23 Electrical Load



Data Center Facility 4 CRAC 7 Operating Conditions



Appendix B

Estimate of Data Center Non-Critical Load

Both Data Center 4.1 and Data Center 4.5 contain a variety of computer equipment. The equipment is categorized as follows.

Table B.1. Computer Equipment Categories

Servers	Racks of modular computers, or unitary systems in cabinets.
Terminals	Monitor-and-keyboard interfaces to the servers.
Tape Drives	Desktop unit for backup tapes.
PCs	Standard combination of monitor, keyboard, and stand-alone computer.
Laser Printers	Desktop models.
Dot-matrix printers	Mostly large floor models; some desktop models.

Table B.2. Computer Equipment Count and Operating Status

	Data Center 4.1		Data Center 4.5	
	Count	On	Count	On
Servers	60	60	66	66
Terminals	6	6	4	4
Tape Drives	1	1	0	0
PCs	15	8	25	17
Laser Printers	1	1	3	3
Dot-matrix	0	0	7	7

All servers, terminals, and tape drives are considered critical load. Approximately half of the PCs are critical load and half are not. The printers are not critical load, but they draw very little power in stand-by mode. None of the printers were operating during the period of monitoring. Printer power is not included in the estimate of non-critical loads.

The PCs draw an estimated 100-Watts each. Assuming four PCs in Data Center 4.1 and eight PCs in Data Center 4.5 are non-critical load, which yields 400 W and 800 W for the data center non-critical loads, respectively.

Appendix C

Estimate of Data-Aire CRAC Unit Performance and Efficiency

All specifications are from Data-Aire's brochure; see www.dataaire.com.

Evaporator Unit

Model No. DAAD-2034

The units in Data Center 4.5 were observed to be equipped with hot gas reheat. Otherwise the units are assumed to be standard configuration.

Parameter	Value	Units	Remarks
Nominal cooling capacity	20	tons	--
Total cooling capacity	247,900	BTU/hr	72 deg F DB, 60 deg F WB, 50% RH
Sensible cooling capacity	201,600	BTU/hr	72 deg F DB, 60 deg F WB, 50% RH
No. of fan motors	1	--	--
Fan motor nameplate power	5	hp	Full Load Amps (FLA) = 6.6 amps
Air flow	8,000	cfm	External static pressure = 0.5 in. w.g.
No. of compressors	2	--	Hermetic scroll
Humidifier power	3.3 to 10.2	kW	--

Rooftop Condenser Unit

Model No. DARC-37

The standard condenser for the DAAD-2034 evaporator unit is the DARC-21, but this facility uses the DARC-37 condenser unit instead.

Parameter	Value	Units	Remarks
No. of fans	4	--	--
Fan motor nameplate power	0.75	hp	Single-phase, permanent split capacitor

Estimates

The Data-Aire brochure does not list EER ratings or other efficiency data. A very similar 20-ton CRAC unit made by Compu-Aire (model CAA-20) has an EER of 8.5. This is a reasonable number; the Data-Aire DAAD-2034 is assumed to have the same efficiency.

EER 8.5 = 1.41 kW/ton.

Assume unit is performing sensible cooling only (no dehumidifying).
Assume unit is not humidifying.

Specifications say the unit's sensible cooling capacity is 201,600 BTU/hr = 16.8 tons.

$16.8 \text{ tons} \times 1.41 \text{ kW/ton} = 23.7 \text{ kW}$ at full load. This electrical draw includes the condenser unit power.

Each fan motor in the condenser unit is nominally 0.75 hp. Assume the brake horsepower is 80% of that, or $0.60 \text{ hp} = 0.80 \text{ kW}$. This agrees well with the measured average condenser power of 0.70 kW, assuming only one fan was running. At full load, assume condenser power is $4 \text{ fans} \times 0.80 \text{ kW/fan} = 3.2 \text{ kW}$.

The CRAC unit fan motor FLA rating is 6.6 amps. Motor full load kW = $6.6 \text{ amps} \times 460 \text{ volts} \times 1.732 = 5.3 \text{ kW}$. Assume the unit is designed so that the motor draws 80% of full load amps, or 4.2 kW. The fan motor is constant speed.

The spot measurement of CRAC 24 power was 4.3 kW. Assume CRAC 24 compressor was not running, only the fan. Assume fans in all CRAC units in Data Center 4.5 draw the same amount of power.

At full load with sensible cooling only, we assume both compressors are running equally. Each compressor draws $(23.7 \text{ kW} - 4.2 \text{ kW} - 3.2 \text{ kW}) / 2 = 8.1 \text{ kW}$.

CRAC 23 drew 12.2 kW during the period of monitoring. Subtracting 4.3 kW for the fan leaves 7.9 kW, which corresponds well with 8.1 kW for a single compressor.

If full load sensible cooling is 16.8 tons with two compressors, then one compressor corresponds to 8.4 tons.

If one other CRAC unit was operating similarly to CRAC 23, then the CRAC units were providing 16.8 tons of cooling to Data Center 4.5.

Appendix D

Marley Technical Report: The Application of Cooling Towers for Free Cooling

This report is available on the World Wide Web at
http://www.marleyct.com/pdf_forms/TRH-002.pdf

**Appendix VI. “Data Center Energy Benchmarking Case Study,
Facility 5”; Federal Data Center Facility, Rumsey Engineers, Inc.;
LBNL - TengFang Xu, 2003**

DATA CENTER ENERGY BENCHMARKING

CASE STUDY, APRIL 2003



DATA CENTER FACILITY 5

Sponsored by



**Lawrence Berkeley National Laboratory
Environmental Energy Technologies Division**

Prepared by



**99 Linden Street
Oakland, CA 94607
(510) 663-2070**

Acknowledgements

Rumsey Engineers and Lawrence Berkeley National Laboratory (LBNL) would thank the anonymous facility managers, engineers, and technicians for their generous assistance and cooperation in this data center benchmarking project. We would also thank Christine Condon of PG&E for providing monitoring equipment. This report benefits from review comments by Steve Greenberg, William Tschudi, and Dale Sartor, and is edited by Tengfang Xu, the Project Manager of LBNL.

Disclaimer

Neither LBNL nor any of its employees makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any data, information, method, product or process disclosed in this document, or represents that its use will not infringe any privately-owned rights, including but not, limited to, patents, trademarks, or copyrights.

Contents

I.	DEFINITIONS	1
II.	INTRODUCTION.....	3
III.	SITE OVERVIEW	3
IV.	ENERGY USE.....	4
	ELECTRICAL EQUIPMENT AND BACKUP POWER SYSTEM.....	4
	COOLING SYSTEM	4
	LIGHTING	6
	SUMMARY OF MEASUREMENTS AND METRICS	6
V.	ENERGY EFFICIENCY RECOMMENDATIONS	9
	TUNE CONTROLS FOR AHU-2.....	9
	INVESTIGATE AHU-5	9
	LIGHTING CONTROL FOR DATA CENTER AND EQUIPMENT ROOMS	10
	RAISE CHILLED WATER TEMPERATURE	10
	FREE COOLING	10
	VARIABLE-PRIMARY ONLY PUMP RETROFIT	11
	COMMISSIONING AND CONTROL IMPROVEMENT	11
APPENDICES		
<i>A -- CHARTS OF MONITORED DATA</i>		

I. Definitions

Chiller Efficiency	The power used (kW), per ton of cooling produced by the chiller.
Measured Computer /Server Load Density	Ratio of actual measured Data Center Server Load in Watts (W) to the gross area (ft ² or sf) of Data Center Floor Space. Includes vacant space in floor area.
Projected Computer /Server Load Density	Ratio of forecasted Data Center Server Load in Watts (W) to square foot area (ft ² or sf) of the Data Center Floor if the Data Center Floor were fully occupied. The Data Center Server Load is inflated by the percentage of currently occupied space.
Computer Load Density – Rack Footprint	Measured Data Center Server Load in Watts (W) divided by the total area that the racks or equivalents occupy, or the rack “footprint”.
Computer Load Density per Rack	Ratio of actual measured Data Center Server Load in Watts (W) per rack. This is the average density per rack.
Cooling Load Density	The amount of cooling (tons) supplied to a given floor space (Ton/ft ² or sf)
Critical Load	Electrical load of equipment that must keep running in the event of a power failure. Such loads are typically served by an Uninterruptible Power Supply (UPS), which uses a bank of batteries to support the load when the normal source of power fails. The batteries can support the load for only a short period. In some facilities the equipment is shut down gracefully and turned off until normal power returns. In other facilities a backup generator, typically diesel-powered, comes on-line and provides power for a longer period of time.
Data Center Cooling	Electrical power devoted to cooling equipment for the Data Center Floor space
Data Center Facility	A facility that contains data storage and processing equipment (servers) associated with a concentration of data cables.
Data Center Floor / Space	Total footprint area of controlled access space devoted to company/customer equipment. Includes aisles, caged space, cooling units, electrical panels, fire suppression equipment, and other support equipment. This gross floor

	space is what is typically used by facility engineers in calculating a computer load density (W/sf). ¹
Computer Equipment Occupancy	This is based on an estimate on how physically loaded the data centers are by computer and equipment footprints. Some people use the term as “Data Center Occupancy.”
Data Center Server/Computer Load	Electrical power devoted to equipment on the Data Center Floor. Typically the power measured upstream of power distribution units or panels. Includes servers, switches, routers, storage equipment, monitors, and other equipment.
Cooling Effectiveness Index	Ratio of electrical power devoted to cooling data space to the electrical power used by computer and equipment in Data Center Floor. A lower number corresponds to more effective cooling.
Cooling Load Tons	A unit used to measure the amount of cooling being done. Equivalent to 12,000 British Thermal Units (BTU) per hour.

¹ Users look at watts per square foot in a different way. With an entire room full of communication and computer equipment, they are not so much concerned with the power density associated with a specific footprint or floor tile, but with larger areas and perhaps even the entire room. Facilities engineers typically take the actual UPS power output consumed by computer hardware and communication equipment in the room being studied (but not including air handlers, lights, etc.) and divide it by the gross floor space in the room. The gross space of a room will typically include a lot of areas not consuming UPS power such as access aisles, white areas where no computer equipment is installed yet, and space for site infrastructure equipment like Power Distribution Units (PDU) and air handlers. The resulting gross watts per square foot (watt/ft²-gross) or gross watts per square meter (watt/m²-gross) will be significantly lower than the watts per footprint measured by a hardware manufacturer in a laboratory setting.

II. Introduction



Rumsey Engineers and the Lawrence Berkeley National Laboratory (LBNL) have teamed up to conduct an energy study as part of LBNL's Data Center Load Characterization under sponsorship by the Federal Energy Management Program (FEMP). Measurements were conducted on-site in October of 2002, with the particular aim of determining the end-use of electricity of data center systems. This study may provide useful information for designers to consider in their decision-making for future design and

construction of data centers. The identity of the organization that owns this data center is kept anonymous. The facility that houses the data center is referred to throughout this report as Data Center Facility 5.

This report will present energy density and the energy efficiency metrics evaluated for the case studies. The variety of mechanical system equipment, and data center types, coupled with the results of the data will provide insight into efficient cooling strategies for data centers. The goal of the project is to obtain benchmarking data and to identify energy efficiency opportunities and best practices.

III. Site Overview

Data Center Facility 5 is located in California's San Francisco Bay Area. The data center occupies 16,000 square feet (sf) of a 62,700 sf multi-level building. The data center houses servers (computer equipment), storage drives and a control room. It operates 24 hours a day with a small number of employees who work with the computers daily. The control room, situated within the data center is 800 sf and is used to troubleshoot and monitor the activity of the computers. The remainder of the building is used as office space.

During this study, the data center was under expansion. Measurements made reflect a data center area of 16,000 sf, which does not include the expansion. According to the building facilities manager, the space allocated to computer equipment in the data center was 60% occupied. The computer equipment and storage drives occupy 18% (2,900 sf) of the data center. The support equipment (i.e. air handlers, electrical power distribution units, and uninterruptible power supply) occupies 31% (5,000 sf) of the data center. The control room occupies 5% (800 sf) of the data center. The remaining 46% (7,300 sf) is either walkway space between the computer and support equipment, or unused floor space.

IV. Energy Use

ELECTRICAL EQUIPMENT AND BACKUP POWER SYSTEM

The requirement for high reliability of this facility was determined not to be critical in considering the associated costs and the best use of data center investment. Therefore, the use of UPS (uninterruptible power supply) systems for the computers in this data center is minimal. There are small UPS systems for the building's phone system and for a bank of computers in the data center. A UPS delivers "clean" power to computer equipment by filtering out voltage spikes and surges. A UPS continuously converts alternating current (AC) power to direct current (DC) power and charges a battery bank. DC power from the batteries is then converted back to AC power to feed the data center equipment. In the event of a power outage the battery bank supplies backup power for a specified time until power is restored or the computer is safely shut down. The UPS system for the computers has a 225-kVA capacity and is currently 10% loaded. Typically, UPS loaded lightly tend to be very poor in efficiency. Most of the computer equipment is served by power distribution units (PDUs) situated throughout the data center. The PDUs receive power from the main electrical distribution boards. Their purpose is to remove spikes and transients and to convert power to 120 Volts from 480 Volts.

COOLING SYSTEM

A chilled water system is utilized to cool the entire building, including the data center space. It is comprised of two Trane 800-ton, water-cooled centrifugal chillers, one of which has a variable frequency drive (VFD). Normally, the VFD-equipped chiller 1 runs with the non-VFD chiller 2 in standby. A third VFD-equipped chiller 3 being added is part of the buildout. The rated nominal efficiency of the VFD chillers 1 and 3 is 0.40 kW/Ton. The rated nominal efficiency of non-VFD chiller 2 is 0.42 kW/Ton.²

There are two, two-cell cooling towers with a 25 hp single-speed fan in each cell, supplied by two constant-speed 75 hp condenser water pumps. Normally, one pump operates at a time on a lead-lag schedule. The chilled-water pumping system employed is a constant-primary with variable-secondary loop. Normally, one 25 hp primary pump and one variable speed 60 hp secondary pump operates. There are identical primary and secondary pumps for backup that operate on a



² Based on 1920 gal per minute (gpm), entering and leaving chilled water temperatures of 55 °F, and 45 °F, respectively, and entering condenser water temperature of 75 °F.

lead-lag schedule. A third set of identical condenser and chilled-water pumps is being added as part of the buildout.

The data center is conditioned by twenty-one computer room air conditioners (CRAC) and a single air handler, AHU-2. An additional 7 computer room AHUs were added as part of the buildout. Another small unit, AHU-4, serves the network room just off the main computer floor. As part of the buildout, the small conference room served by AHU-3 is being merged with the network room. Air handler AHU-2 is used to supply outside air for ventilation requirements and to positively pressurize the data center. The CRAC units are supplied with chilled water and have variable speed drives on the fans that deliver the cooled air through the 3-foot raised floor. The CRAC units have been carefully distributed within the data center to maximize the delivery of conditioned air. They are not all situated along the perimeter of the data center as typically designed for data centers. Simply placing CRAC units along the perimeter of a data center could cause short-circuiting of the conditioned airflow, and result in warmer temperatures towards the inner space of a data center. In addition, the 3-foot raised floor allows for low air pressure drop when supplying the same airflow rate, hence lowers fan power to deliver the same airflow within a certain time frame.

Electrical consumption data from chilled water plant and the data center's HVAC equipment is listed in the table below. Measurements were made on all equipment served with chilled water; either with data collected through the automated control system or with field attached equipment. Please refer to the Appendix for graphs of the measurements over the entire monitored period.

TABLE 1. COOLING EQUIPMENT ELECTRICAL AND LOAD MEASUREMENTS

Equipment	Spot / Calculated	Date	Units	Measurement
Chiller Power	Spot	10/25/02	kW	118
Chiller Tonnage	Spot	10/25/02	Ton	293
Pump Power	Spot	10/25/02	kW	94
CRAC Unit Power	Spot	10/25/02	kW	61
AHU-2 Power	Spot	10/25/02	kW	0.25

It was necessary to identify the load solely to the data center, in order to segregate the chilled water plant power consumption attributed to cooling the data center (please see the following table).

TABLE 2. DATA CENTER COOLING EQUIPMENT ELECTRICAL AND LOAD MEASUREMENTS

Equipment	Spot / Calculated	Date	Units	Measurement
Chiller Tonnage attributed to Data Center	Calculated	10/25/02	Ton	256.3
CRAC Unit Tonnage	Calculated	10/25/02	Ton	255.6
AHU-2 Tonnage	Calculated	10/25/02	Ton	0.6
Chiller Power attributed to Data Center	Calculated	10/25/02	kW	103
Cooling Tower Fan Power attributed to Data Center	Calculated	10/25/02	kW	9
Pump Power attributed to Data Center	Calculated	10/25/02	kW	82

The load that the CRAC units and AHU-2 delivered to the data center was 87% of the chiller's entire load. This stipulates that a majority of the heat load within the building is generated from the data center. On October 25, 2002, the outside air temperature during the day was 60 °F on average. Due to the mild outdoor air conditions, the office spaces were demanding a small load of only 37 tons from the chilled water, the remainder being supplied by the outside-air economizer.

LIGHTING

Lighting in the data center consists of ceiling-mounted, troffer-type fluorescent fixtures. All lights were fully on when taking a spot power measurement. The spot measurement of the data center lighting power indicated approximately 16 kW when all lights are on. The actual lighting power may change over time because there are different schedules.


SUMMARY OF MEASUREMENTS AND METRICS

The table below shows the exemplar power consumption of HVAC and lighting equipment in the data center, which is also shown graphically in the pie chart below.

TABLE 3. DATA CENTER ELECTRICAL POWER USAGE

Computer Loads	565 kW	68%
HVAC - Air Movement	61 kW	7%
HVAC - Chiller Plant	194 kW	23%
Lighting	16 kW	2%

FIGURE 1. DATA CENTER ELECTRICAL CONSUMPTION



The computer loads amount to 68% of the data center power usage. Pumping and cooling of the chiller systems is the second largest consumer at 23%, with air movement at 7% as the next largest. Together, the HVAC components amount to 30% of the data center's power usage. Lighting accounts only 2% of the total power usage.



The electrical and cooling loads can be represented by different metrics. The most commonly used metric among data center facilities is the computer load density in watts consumed per square foot. However, this square footage used is not consistent between designers and is a source of problems.³ The Uptime Institute's definition of "Data Center Floor Area" includes the entire area that is dedicated to data center equipment. This

³ See "Data Center Power Requirements: Measurements from Silicon Valley", by Mitchell-Jackson, Koomey, Nordman, & Blazek, December 2001. It is available on the web at http://enduse.lbl.gov/Info/Data_Center_Journal_Article2.pdf.)

includes rack spaces, storage areas, aisles, and areas utilized for power distribution and CRAC units. This definition is however basically subjective. Some data centers use kVA/rack or kW/rack as a design parameter. We have also calculated the W/sf based on the rack area alone. In addition to the previous metrics, the “non-computer” power densities are calculated, based on the “data center floor area”. The “non-computer” items include the equipment such as HVAC, and lighting. In addition, based upon actual 60% computer and equipment occupancy in the data center, the projected computer load density at otherwise 100% occupancy can be estimated.

TABLE 4. DATA CENTER ELECTRICAL CONSUMPTION METRICS

Data Center Floor Area	16,000	sf
Rack Area	2,880	sf
Data Center Occupancy	60	%
<i>Based on Data Center Floor Area:</i>		
Computer Load Density	35	W/sf
Non-Computer Load Density	17	W/sf
Projected Computer Load Density	59	W/sf
<i>Based on Rack Area:</i>		
Computer Load Density	196	W/sf
Projected Computer Load Density	327	W/sf

The computer load density based on the “data center floor area” (gross area) is 35 W/sf. At full occupancy, the computer load density is projected to be 59 W/sf. The computer load density based on “rack area” is presently 196 W/sf, and is projected to be 327 W/sf at full occupancy. The non-computer load (or power) density, which includes HVAC and lighting is 17 W/sf.

Since the loading of data centers and computer types are site specific, a more useful metric for evaluating how efficiently the data center is cooled can be represented as a ratio of cooling power to computer power. The “Cooling Effectiveness Index ” is 0.45 kW/kW (a lower number corresponds to more effective cooling, see Table 5). This indicates that in Data Center 5 for every unit of power used by computer equipment, 45% of the cooling power is being called for to maintain the conditions of the data center.

Although there is a small amount of cooling load generated from the presence of human activities, the load is insignificant compared to the computer loads. The more traditional metrics of power per ton of cooling (kW/Ton) are calculated for total HVAC efficiency (chillers, pumps, and air handlers), and for the chillers.

TABLE 5. HVAC EFFICIENCY METRICS

	Value	Units
Cooling kW: Computer Load kW	0.45	--
Cooling Provided by Chilled Water System	293	Tons
Chiller 2 Efficiency	0.4	kW/ton
Chiller 2 Design Efficiency	0.4	kW/ton

Chiller ARU-2 was operating close to the full-load efficiency it was designed for. Constant speed chillers typically operate at poor efficiencies at part load. However, the measured efficiency is positively impacted by a lower condenser water temperature which is lower than the designed water temperature.

V. Energy Efficiency Recommendations

As a result of this study, the team was able to identify energy efficiency opportunities for the facility. The trended data in particular reveals the operation of the mechanical equipment and if the equipment is operating as intended.

TUNE CONTROLS FOR AHU-2

The air handler serving the data center, AHU-2 is excessively cycling between heating and cooling modes. A figure of the unit's temperatures and valve positions can be seen in the Appendix. Tuning of the controls is recommended for the air handler to minimize the cycling between heating and cooling.

INVESTIGATE AHU-5

The air handler serving the telephone equipment room, AHU-5 is overcooling the space. The set point of the space was designed for 75°F. During the monitoring period, the unit was continuously running with an average return air temperature of 64 °F.⁴ This implies that the room temperature is 64°F. The potential culprit of this wasted energy is likely to be the room temperature sensor or a control problem.

⁴ Return air temperature was measured at the return air path of the air handling unit, AHU-5.

LIGHTING CONTROL FOR DATA CENTER AND EQUIPMENT ROOMS

The data center is sparsely occupied during business hours. It has even lower or no occupancy after normal business hours.

In each of the two largest basement equipment rooms, a portion of the lights is currently controlled with an on/off switch while the remainder stays on continuously. Lights in the equipment rooms are typically left 100% on by the occupant(s) of the rooms since on/off toggle switches are utilized. Adding timer switches for the controlled lighting was a recommendation made by the mechanical systems technician. By replacing the switches with timer switches, lighting power consumption can be reduced in this 24-hour operated facility. An occupancy sensor switch is an alternative, however, multiple sensors may be desired due to the obstructions in the rooms. Installing multiple overhead occupancy sensors for lighting control would ensure that the lights remain on when a person is working with the computer equipment.

RAISE CHILLED WATER TEMPERATURE

Other operating conditions remaining the same, a chiller that produces chilled water with warmer water temperature would normally operate more efficiently. It is possible that the chilled water (CHW) temperature can be raised from the current 45°F to somewhere in the range of 50°F to 55°F. During the study, the chiller was only 37% loaded, indicating that the control valves on the CRAC and air handler's cooling coils were only partially open. By raising the CHW temperatures, the valves should open more fully to allow more CHW flow through the coils. In addition, energy savings can be achieved by improving chilled water temperature control and improving performance monitoring and tracking.

The dehumidifying control in the data center has been disabled, which indicates there probably is no practical need for chilled water temperatures as low as 45 °F. However, humidity requirements for the office spaces should also be evaluated.

FREE COOLING

Free cooling can be achieved by displacing the chillers and using the cooling towers to produce water around 50°F when outdoor wet bulb temperatures are relatively low, e.g., below 45°F. This measure can be directly related to the recommendation above, "Raise Chilled Water Temperature." A retrofit would be necessary to implement free cooling. During the year, free cooling can be utilized for 1,260 hours. For a chiller running at an average of 150 kW for 1,260 hours, 189,000 kWh will be consumed. Significant energy savings could be achieved when using the low-energy free cooling strategy. Consider water side economizer so that more efficiency could be attained. The economics of implementing a free cooling system should however be examined carefully. In addition, improve cooling tower controls by using variable or multi-speed motors on tower fans and pumps.

VARIABLE-PRIMARY ONLY PUMP RETROFIT

The chilled water distribution is configured in a constant-primary with variable-secondary loop. The cumulative pumping power was determined to be 94 kW during the monitoring period. This was equal to 80% of the chiller power at 118 kW. Retrofitting the current configuration with a variable-primary only pumping system may create pumping energy savings. The existing pressure transducer used to control the secondary pump speed can be reused to implement the variable-primary only pumping system.


It is worthwhile to investigate the current pumping set points and control in AHUs to avoid unwanted bypass.

COMMISSIONING AND CONTROL IMPROVEMENT


The improvement in commissioning, monitoring sensors (e.g., regular calibrations and additions where desired), re-heat control in AHU systems, and airflow control would add to the efficiency gains. To improve air management such as efficient VFD motors of CRACs and optimized distribution of CRAC units also helps to increase efficiency.

Appendix A – Monitored Data


Data Center Facility 5 OSA Conditions




Data Center Facility 5
Transformer 529A Power (Office Plug Loads)




Data Center Facility 5
Panel 526A2A Power (Equipment Room and Misc Lighting)




Data Center Facility 5 Primary Chilled Water Temperatures




Data Center Facility 5 AHU-2 Conditions



Data Center Facility 5
AHU-2 Mixed Air Temp



Data Center Facility 5 AHU-5 Air Temperatures



**Appendix VII. “Data Center Energy Benchmarking Case Study,
Facility 6”; PIER Project, Rumsey Engineers, Inc.; LBNL –
William Tschudi, 2003**

DATA CENTER ENERGY BENCHMARKING CASE STUDY

DECEMBER 2002



FACILITY 6

SPONSORED BY:



LAWRENCE BERKELEY NATIONAL LABORATORY

PREPARED BY:



**99 LINDEN STREET
Oakland, CA 94607
(510) 663-2070**

Acknowledgements

Rumsey Engineers is grateful to the facility managers/directors, engineers and technicians for their generous assistance and cooperation. Special thanks to Christine Condon of PG&E for providing monitoring equipment on short notice. Thanks to the Lawrence Berkeley National Laboratory (LBNL), and the California Energy Commission for funding this project.

Disclaimer

Neither Rumsey Engineers, LBNL nor any of its employees makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any data, information, method, product or process disclosed in this document, or represents that its use will not infringe any privately-owned rights, including but not, limited to, patents, trademarks, or copyrights.

Contents

I.	EXECUTIVE SUMMARY.....	1
II.	DEFINITIONS	6
III.	INTRODUCTION.....	8
IV.	SITE OVERVIEW	8
V.	ENERGY USE – DATA CENTER 6.1.....	9
	DATA CENTER 6.1: ELECTRICAL EQUIPMENT AND BACKUP POWER SYSTEM.....	9
	DATA CENTER 6.1: COOLING SYSTEM.....	10
	DATA CENTER 6.1: LIGHTING	13
	DATA CENTER 6.1: SUMMARY MEASUREMENTS AND METRICS	13
VI.	ENERGY USE – DATA CENTER 6.2	18
	DATA CENTER 6.2: ELECTRICAL EQUIPMENT AND BACKUP POWER SYSTEM.....	18
	DATA CENTER 6.2: COOLING SYSTEM.....	18
	DATA CENTER 6.2: LIGHTING	20
	DATA CENTER 6.2: SUMMARY MEASUREMENTS AND METRICS	20
VII.	ENERGY EFFICIENCY RECOMMENDATIONS	24
	GENERAL GUIDELINES FOR FUTURE CONSTRUCTION	24
	<i>Efficient Chilled Water System</i>	24
	<i>Air Management</i>	26
	<i>Air Management – Rack Configuration</i>	26
	<i>Commissioning of New Systems and Optimized Control Strategies</i>	26
	<i>Lighting Controls</i>	27
	DATA CENTER 6.1 SPECIFIC OBSERVATIONS.....	27
	DATA CENTER 6.2 SPECIFIC OBSERVATIONS.....	28

APPENDICES

CHARTS OF MONITORED DATA

I. Executive Summary

Rumsey Engineers and the Lawrence Berkeley National Laboratory (LBNL) have teamed up to conduct an energy study as part of LBNL's Data Center Load Characterization and Roadmap Project, under sponsorship by the California Energy Commission (CEC). This study is intended to provide measured information on energy and power use in data centers, and to help designers make better decisions about the design and construction of data centers in the near future. This report describes the outcomes of energy benchmarking in two data centers in Northern California, and the observations on potential opportunities in efficiency improvement. Data centers at different organizations in Northern California were analyzed, with the particular aim of determining the end-use of electricity.

This report documents the findings for one of the case studies – termed Data Center Facility 6. Additional case studies and benchmark results as they become available will be provided on LBNL's website (<http://datacenters.lbl.gov>) For comparison purposes, the results of a similar benchmarking study completed for the Pacific Gas and Electric Company (PG&E) in 2001 are included in this report.

Facility 6 contains two data centers, in two separate office buildings. These data centers contain mainly server type computers and data storage devices and resemble the server farms that became common as a result of the Internet Age.¹ Both Data Center 6.1 and Data Center 6.2 areas in facility A represent approximately 3% of the total building area. This percentage is a relatively small percentage, therefore the end use electricity of the whole building was not evaluated. Data Centers 6.1 and 6.2 were each 2,400, and 2,500 square feet (sf) respectively. Both data centers were primarily cooled by chilled water feeding computer room air handlers (Data Center 6.1), or fan coil units (Data Center 6.2). Both data centers were conditioned with overhead supply air and did not utilize raised floors.


The current computer energy loads are listed in the table below. A qualitative estimate of the loading of the racks was made, and the future computer energy loads were estimated based on this loading. For comparison purposes the computer loads of the data centers studied in the PG&E project are also included (Data Centers 1, 2, and 3). The computer loads are also shown graphically.


The measured computer load densities at Facility 6 are greater than the computer load densities measured in the previous PG&E study. The measurements project full occupancy densities of 81 and 95 W/sf, which are considerably higher than the full occupancy density projected in the PG&E study. The remaining energy loads of Data Centers 6.1 and 6.2 include air conditioning loads, lighting, and uninterruptible power supply inefficiencies. They are shown in graphical format below, as well as tabular format in the report.

¹ Based on the rack configuration, high density of computers, and absence of the large mainframe servers that were common in older data centers.

CURRENT AND FUTURE COMPUTER LOADS

Data Center	Data Center Area (sf)	Computer Load (kW)	Computer Load Energy Density (W/sf)	Occupancy (%)	Projected Computer Load Energy Density (W/sf)	Number of Racks	kW/Rack
6.1	2,400	155	65	80%	81	101	1.5
6.2	2,500	119	48	50%	95	83	1.4
1	62,870	1,500	24	75%	32	--	--
2	60,400	2,040	34	65%	52	--	--
3	25,000	1,110	44	85%	52	--	--





A large percentage, approximately 73%, of the total electrical load is from the computer loads. However, the HVAC loads totaled 21%. Since this represents a large percentage, efficiency improvements could result in significant energy savings. A number of energy efficiency metrics were calculated, including UPS efficiency, chiller efficiency, and air handler efficiency. A useful efficiency metric, particularly for data centers is a cooling efficiency, calculated as a ratio of cooling energy to computer load energy. These are shown below.


A more detailed discussion is presented in the report. In summary, the measured efficiencies of the chillers were approximately equal to their design efficiencies, as would be expected for the operating conditions. This is because the design efficiencies are based on 95 °F entering condenser temperature. When outdoor air temperatures are below this temperature, the chiller can reject energy more easily, and therefore has lower power consumption. Based on the outdoor air conditions in this geographical area, better efficiencies are expected. The air handler efficiencies were below their design efficiencies; this is likely due to excess pressure losses through the ducting.

DATA CENTER 6.1 EFFICIENCY METRICS

Efficiency Metric	Value	Units
UPS Efficiency	94%	--
Cooling kW: Computer Load kW	0.3	--

Efficiency Metric	Value	Units
Average Chiller 1 (40 Ton) Efficiency	0.9	kW/Ton
Average Chiller 2 (100 Ton) Operating Efficiency	1.0	kW/Ton
Chiller 1 Design Efficiency	1.1	kW/Ton
Chiller 2 Design Efficiency	1.3	kW/Ton
AHU 1 Efficiency – Measured	1,367	CFM/kW
AHU 2 Efficiency – Measured	1,375	CFM/kW
AHU 3 Efficiency – Measured	1,387	CFM/kW
AHU 1 Design Efficiency	2,221	CFM/kW
AHU 2 Design Efficiency	2,044	CFM/kW
AHU 3 Design Efficiency	3,219	CFM/kW
Overall HVAC Efficiency	1.3	kW/Ton

The electrical energy characteristics for Data Center 6.2 are shown in the graph below.



In this case, the HVAC loads, at 35%, represented an even larger percentage of the total energy use. Similar opportunities for energy savings exist, and are described in detail in the report. The energy efficiency metrics are listed in the table below. Since Data Center 6.2 utilized fan coil units, rather than computer room air handlers, it was not practical to obtain individual efficiencies.

DATA CENTER 6.2 EFFICIENCY METRICS

Efficiency Metric	Value	Units
UPS1 Efficiency	93%	%
UPS2 Efficiency	90%	%
Cooling kW: Computer Load kW	0.6	--
Chiller 1 Efficiency	1.0	kW/ton
Chiller 2 Efficiency	1.1	kW/ton
Chiller 1,2 Design Efficiency	1.3	kW/ton
Fan Coil Unit Design Efficiency	2370	CFM/kW
Overall HVAC Efficiency	1.6	kW/ton

The chiller efficiency results were comparable to the efficiencies of chillers serving Data Center 6.1. Again, the overall efficiencies are low, as would be expected from air-cooled chillers. Though the efficiencies are comparable to the design efficiencies, better performance is expected, since the operating conditions are more favorable, as discussed earlier. The design efficiencies of the FCUs are comparable to the design efficiencies of the AHUs used in Data Center 6.1, though the actual efficiencies were not measured.

II. Definitions

Data Center Facility	A facility that contains both central communications equipment, and data storage and processing equipment (servers) associated with a concentration of data cables. Can be used interchangeably with Server Farm Facility
Server Farm Facility	A facility that contains both central communications equipment, and data storage and processing equipment associated with a concentration of data cables. Can be used interchangeably with Data Center Facility. Also defined as a common physical space on the Data Center Floor where server equipment is located (i.e. server farm)
Data Center Floor / Space	Total footprint area of controlled access space devoted to company/customer equipment. Includes aisle ways, caged space, cooling units, electrical panels, fire suppression equipment, and other support equipment. Per the Uptime Institute Definitions, this gross floor space is what is typically used by facility engineers in calculating a computer load density (W/sf). ²
Data Center Occupancy	This is based on a qualitative estimate on how physically loaded the data centers are.
Data Center Cooling	Electrical power devoted to cooling equipment for the Data Center Floor space
Data Center Server/Computer Load	Electrical power devoted to equipment on the Data Center Floor. Typically the power measured upstream of power distribution units or panels. Includes servers, switches, routers, storage equipment, monitors, and other equipment.
Computer/Server Load Measured Energy Density	Ratio of actual measured Data Center Server Load in Watts (W) to the square foot area (ft ² or sf) of Data Center Floor. Includes vacant space in floor area

² Users look at watts per square foot in a different way. With an entire room full of communication and computer equipment, they are not so much concerned with the power density associated with a specific footprint or floor tile, but with larger areas and perhaps even the entire room. Facilities engineers typically take the actual UPS power output consumed by computer hardware and communication equipment in the room being studied (but not including air handlers, lights, etc.) and divide it by the gross floor space in the room. The gross space of a room will typically include a lot of areas not consuming UPS power such as access aisles, white areas where no computer equipment is installed yet, and space for site infrastructure equipment like Power Distribution Units (PDU) and air handlers. The resulting gross watts per square foot (watt/ft²-gross) or gross watts per square meter (watt/m²-gross) will be significantly lower than the watts per footprint measured by a hardware manufacturer in a laboratory setting.

Computer Load Density – Rack Footprint	Measured Data Center Server Load in Watts (W) divided by the total area that the racks occupy, or the rack “footprint”.
Computer Load Density per Rack	Ratio of actual measured Data Center Server Load in Watts (W) per rack. This is the average density per rack.
Computer /Server Load Projected Energy Density	Ratio of forecasted Data Center Server Load in Watts (W) to square foot area (ft ² or sf) of the Data Center Floor if the Data Center Floor were fully occupied. The Data Center Server Load is inflated by the percentage of currently occupied space.
Cooling Load Tons	A unit used to measure the amount of cooling being done. Equivalent to 12,000 British Thermal Units (BTU) per hour.
Chiller Efficiency	The power used (kW), per ton of cooling produced by the chiller.
Air Handler Efficiency 1	The air flow (CFM) per power used (kW) by the CRAC unit fan
Air Handler Efficiency 2	The power used (kW), per ton of cooling achieved (ton) by the air-handling unit.
Cooling Load Density	The amount of cooling (tons) in a given area (ft ² or sf)
Air Flow Density	The air flow (CFM) in a given area (ft ² or sf)

III. Introduction

This report describes the measurement methodology and results obtained for this case study. The facility contained two Data Centers, which were measured independently. In each data center, electricity end use was determined. This means that the energy consumed by all equipment related to the data center was measured. Such equipment includes the actual computer power consumption, the data center air conditioning equipment, the lighting, and the inefficiencies associated with the uninterruptible power supply (UPS). The computer load density is also determined based on the gross area of the data center. This number, in watts per square foot (W/sf) is the metric typically used by facility engineers to represent the power density. Based on a qualitative observation of the data center occupancy, the computer load density at full occupancy is extrapolated. In addition to the typical W/sf metric, the density is also calculated based on the number of racks, and the rack footprint.

Additional information was collected so that the efficiencies of the cooling equipment could be calculated. These efficiencies are compared to the design efficiencies. Opportunities for energy efficiency improvements are described, which are based on observation of the mechanical system design, and measured performance. General design guidance is presented for consideration in future construction. Data Center specific recommendations are made for the as-built systems.

IV. Site Overview

Facility 6 is located in Silicon Valley in California. Two data centers were monitored for energy consumption at Facility 6. The data centers are in separate office buildings, and constitute a relatively small percentage of the total building area. (less than 10%) The data centers, hereafter referred to as Data Center 6.1, and Data Center 6.2, are 2,400 square feet (sf), and 2,500 sf, respectively. Since the data centers represent a small percentage of the overall building area, whole building power consumption is not relevant to determining the *data center* power consumption, and was not monitored. Both data centers house servers and storage drives, and operate 24 hours a day. One of the data

centers serves corporate needs (Data Center 6.1), while the other is mainly used for research and development of new engineering products (Data Center 6.2). Occasionally, during normal business hours, a small number of employees may be in the data centers working with the computers.



Figure – Computer Servers

V. Energy Use – Data Center 6.1

DATA CENTER 6.1: ELECTRICAL EQUIPMENT AND BACKUP POWER SYSTEM

The facility utilizes a Balanced Power 225 kVA uninterruptible power supply (UPS) to provide a constant supply of power to the data center at constant delivery voltage (480/277 V). The UPS converts AC current and stores it as DC current in multiple battery packs. When the voltage is needed, it is converted back to AC current. In the event of a power loss, a 400 kW diesel generator will provide power for approximately 10 hours.

Spot power measurements were taken at the UPS, both at the input and output in order to determine computer plug loads, as well as losses at the UPS system.

TABLE 1. UPS ELECTRICAL MEASUREMENTS

	Electrical Use ¹	Units
UPS Input	164.7	kW
UPS Output	154.9	kW
UPS Losses	9.8	kW
UPS Efficiency	94.0	%

¹ Average measurement taken on 8/21/02, using the PowerSight Power Meter.

DATA CENTER 6.1: COOLING SYSTEM

The data center is cooled separately from the remainder of the building. A chilled water system cools the data center, as well as several small computer labs. It consists of two Trane air-cooled chillers, a 40 Ton scroll chiller, and a 100 Ton rotary chiller. The nominal efficiencies of the chillers are 1.1 and 1.3 kW/Ton, respectively.³ The 100-ton chiller is served by the emergency distribution panel (EDP), and is the primary chiller, though the 40 Ton chiller is often run in unison to ensure sufficient supply of chilled water. The chilled water pumps are 1.5 hp (hydraulic horsepower; brake horsepower unlisted) pumps, and are variable speed, controlling based upon a differential pressure set point. A controlled bypass ensures minimum flow through the chillers. The chilled water system branches off into two feeds, one that is dedicated to the data center, and the other that feeds the computer labs.

Power consumption, flow, and chilled water temperatures⁴ were measured at each chiller over a period of several days. This was to determine the chiller efficiency over a period of varying temperatures.

The computer room air conditioners are constant-speed air handler units (AHU) that are supplied chilled water. There are three air handlers in total, with total cooling capacities of 286,900 British thermal units per hour (Btu/hr), 551,700 Btu/hr, and 185,400 Btu/hr and design airflows of 9200 cubic feet per minute (cfm), 12,700 cfm, and 8,000 cfm, respectively.⁵ Air is returned through grills in the front of the AHU, and exits from the top to ducts that feed the ceiling diffusers. The computer room air handlers control the return air temperature of 70 °F. In addition to the air that is recirculated and cooled by the computer room air handlers, ventilation air is supplied by the main building air conditioning unit. The air handlers do not have humidity control.

Spot measurements of flow, and temperatures were performed at the AHU chilled water supply lines.⁶ In addition, flow rate, supply and return chilled water temperatures to all three handlers were monitored over a period of several days.⁷ It was necessary to identify the chilled water supplied solely to the data center, in order to segregate the chiller power consumption due to cooling of the data center only. Spot measurements of airflow

³ Converted from the EER listed on the equipment schedules. The schedule for the 100-ton chiller was incomplete, and therefore, its EER was assumed to be the same as the identical model chillers that are installed for Data Center 6.2. The nominal loads are based on entering evaporator water temperature of 56 °F, leaving evaporator water temperature of 44 °F, entering condenser air temperature of 95 °F, and flow rates of 80 gpm, and 200 gpm.

⁴ These were measured using an Elite power measuring instrument, an ultrasonic flow meter for pipe flow, and thermistors inserted in the Pete's plugs at the inlet and outlet of the chilled water line.

⁵ The numbering refers to the numbering physically on the units. (CRU #1, CRU #2, CRU #3). This does not correspond with the numbering on the equipment schedule, based on the anticipated motor kW.

⁶ These measurements were taken by measuring pressure drop across the circuit setter on the chilled water line, and by measuring temperatures at Pete's Plugs on the supply and return lines.

⁷ These measurements were made at the main branch that feeds only these units. Measurements of chilled water temperatures were performed by inserting thermistor probes between insulation and the pipe surface. Flow measurements were made using an ultrasonic flow meter.

through the AHUs were measured along with the AHU power consumption to determine how efficiently air is moved.⁸

The spot measurements, and average of trended measurements are listed in the table below. Please refer to the Appendix for graphs of the measurements over the entire monitored period. The chiller pump and chiller power are proportioned to the data center-cooling load in order to properly determine electrical end use for the data center.



Figure – Data Center Air Handling Unit

⁸ Airflow was taken by multiplying the average velocity across the return grille with the grille area, where the velocity was taken with a Shortridge velocity grid.

TABLE 2. COOLING EQUIPMENT ELECTRICAL AND LOAD MEASUREMENTS

Equipment	Spot / Monitored	Date	Measurement	Units
Chiller Pumps - Total	Spot	8/21/02	4.0	kW
Chiller Pumps - Proportioned based on Data Center Load	Spot	8/21/02	2.0	kW
AHU 1 (Compuaire C)⁹	Spot	8/21/02	3.7	kW
AHU 2 (Compuaire B)¹⁰	Spot	8/21/02	4.7	kW
AHU 3 (Compuaire A)¹¹	Spot	8/21/02	1.8	kW
AHU 1 Tonnage	Spot	8/21/02	12	Tons
AHU 2 Tonnage	Spot	8/21/02	16	Tons
AHU 3 Tonnage	Spot	8/21/02	7	Tons
AHU 1 Airflow	Spot	9/4/02	5,086	CFM
AHU 2 Airflow	Spot	9/4/02	6,494	CFM
AHU 3 Airflow	Spot	9/4/02	2,432	CFM
DC Cooling Load From Chilled Water - Based on AHU Tonnage	Spot	8/21/02	124.0	kW
DC Cooling Load From Chilled Water - From Monitoring of Chilled Water Use	Monitored	8/30/2002 - 9/4/2002	111.0	kW
Chiller 2 Total (100 ton)	Spot	8/21/02	48.0	kW
Chiller 1 (40 ton)	Spot	8/21/02	16.0	kW
DC Chiller kW From Spots 1	Spot	8/21/02	35.4	kW
DC Chiller kW From Monitoring - Average	Monitored	8/30/2002 - 9/4/2002	32.3	kW

1 Individual chiller kW proportioned based on the data center cooling load versus total chiller load. This value will vary even if the data center load stays constant, when the chiller load changes, as the efficiency of the chiller is not constant.

⁹ Supply Fan Schedule: 9200 cfm, 5 BHP.

¹⁰ Supply Fan Schedule: 12700 cfm, 7.5 BHP.

¹¹ Supply Fan Schedule: 8000 cfm, 3 BHP.

DATA CENTER 6.1: LIGHTING

Lighting in the data center consists of T-8 tubular fluorescent lamps, and all lights were on when taking power measurements. Lighting Power: 1.16 kW (Taken on 8/21/02) or 0.5 W/sf.


DATA CENTER 6.1: SUMMARY MEASUREMENTS AND METRICS

The table below summarizes the equipment electrical measurements for the data center.

TABLE 3. SUMMARY OF ELECTRICAL MEASUREMENTS

Computer Loads	154.9	kW	73%
UPS Losses	9.8	kW	5%
HVAC - Air Movement	10.0	kW	5%
HVAC - Pumps and Chiller	34.0	kW	16%
Lighting	1.1	kW	1%
Total Energy Use	210.0	kW	100%

These results are shown graphically in the pie chart below.



The computer loads, based on the UPS power supply amounts to 73% of the data center energy usage. Pumping and cooling energy is the second largest consumer at 16%, and air movement specifically is 5%. Together, the HVAC component amounts to 21% of data center energy use, a very significant amount. Therefore, efficiency improvements in energy for HVAC could be significant. Losses in the UPS account for 5% of the data center energy consumption. These losses are more than the lighting, which amounted to only 1% of total energy use.

The electrical and cooling loads can be represented by different metrics. The most commonly used metric among mission critical facilities is the computer load density in watts consumed per square foot. However, the square footage is not always consistent between designers. Some data centers use kVA/rack or kW/rack as a design parameter. Our definition of “Data Center Floor Area” includes the gross area of the data center, which includes rack spaces, aisle spaces, and areas that may eventually contain computer equipment. Per the Uptime Institute, the resulting computer load density in watts per square foot is consistent with what facility engineers use, though this is different from the “footprint” energy density that manufacturers use. We have also calculated the W/sf based on the rack area alone. In addition to the previous metrics, the “non-computer” energy densities are calculated, based on the “data center area”. Using the data center occupancy¹² the computer load density at 100% occupancy is projected.

¹² A qualitative assessment of how physically full the data center is. In this facility, occupancy was determined by a visual inspection of how full the racks in place were.

TABLE 4. ELECTRICAL CONSUMPTION METRICS

Data Center Gross Area ¹	2,400	sf
Rack Area	630	sf (Calculated from a total of 121 racks, and area of 1 rack)
"Occupied" %	80%	Estimated from visual inspection.
Based on Gross Area:		
Computer Load Density	65	W/sf
Non-Computer Load Energy Density	23	W/sf
Projected Computer Load Density	81	W/sf
Based on Rack Area: ²		
Computer Load Density	246	W/sf
Projected Computer Load Density	307	W/sf
On an Individual Rack Basis: ³		
Computer Load Density	1.3	kW/Rack
Projected Computer Load Density	1.6	kW/Rack

1 Gross area includes spaces between racks; does not include entire building area.

2 This is an important metric, because the data center gross area can vary depending on spacing between racks.

3 This is the average rack computer load.

The computer load density based on the data center area (gross area) is 65 W/sf. At full occupancy, the computer load density is projected to be 81 W/sf. The computer load density based on *rack area* is presently 246 W/sf, and is projected to be 307 W/sf at full occupancy. The average computer load, based on the number of racks is currently 1.3 kW/Rack, projected to be 1.6 kW/Rack at full capacity. The non-computer energy density, which includes HVAC, lighting, and UPS losses, is measured at 23 W/sf.

Since the rack density within data centers and computer types are site specific, a more useful metric for evaluating how efficiently the data center is cooled can be represented as a ratio of cooling power to computer power. The “theoretical cooling load” is the same

as the sum of the computer loads and lighting loads, together being the plug loads. (There is a small amount of human activity; however, the energy load is insignificant compared to the computer loads.) This is a good cross check of measurements and may also be an indication of the level of cooling that is provided by non data-center dedicated cooling equipment (i.e., general office building, or “house” air to achieve minimum ventilation). The more traditional metrics of energy per ton of cooling (kW/Ton) are calculated for total HVAC efficiency (chillers, pumps, and air handlers), and for the chillers. The air handler efficiency is based on how much air is actually being moved for the measured power consumption.

TABLE 5. HVAC EFFICIENCY METRICS

Metric	Value	Units
Cooling kW: Computer Load kW	0.3	--
Theoretical Cooling Load *	47	Tons
Cooling Provided by AHUs and Chilled Water	32	Tons
Cooling Provided by House Air (Based on Energy Balance)	13	Tons
Combined Chiller Efficiency	1.0	kW/Ton
Average Chiller 1 (40 Ton) Efficiency	0.9	kW/Ton
Chiller 1 Design Efficiency ¹³	1.1	kW/Ton
Average Chiller 2 (100 Ton) Operating Efficiency	1.0	kW/Ton
Chiller 2 Design Efficiency ¹⁴	1.3	kW/Ton
Overall HVAC Efficiency	1.3	kW/Ton
AHU 1 Efficiency – Measured	1,367	CFM/kW
AHU 2 Efficiency – Measured	1,375	CFM/kW
AHU 3 Efficiency – Measured	1,387	CFM/kW
AHU 1 Design Efficiency ¹⁵	2,221	CFM/kW
AHU 2 Design Efficiency	2,044	CFM/kW
AHU 3 Design Efficiency	3,219	CFM/kW

* Based on computer loads, lighting loads, and fan energy.

¹³ The nominal efficiencies cannot be directly compared to the average operating efficiencies, since the nominal efficiencies are based on full load capacities, and the specific conditions cited previously.

¹⁴ Same as above.

¹⁵ The fan kW is calculated using the schedule fan Bhp and an assumed motor efficiency of 90%. Also, please note the numbering has been changed from the equipment schedule to match the numbering on the units.

From the above table it is shown that the “cooling efficiency” is 0.3 kW/kW. This, however, is based on a cooling load that is below the theoretical cooling load by 30%. This suggests that significant cooling is being achieved by the whole building cooling system (package units). The efficiency and operation of this system was not evaluated. However, the whole building system has the ability to provide cooling by supplying outdoor air when the weather is favorable (i.e., economizing), a very efficient way of providing cooling.

The average chiller efficiencies are slightly better than the design efficiencies, which are at ARI conditions. This is expected since the ARI conditions assume 95°F entering condenser air temperature, which is higher than the average temperatures experienced during the monitored period. When outdoor air temperatures are below this temperature, the chiller can reject energy more easily, and therefore has lower power consumption. Based on the outdoor air conditions in this area, better efficiencies are expected. For every 1 °F drop in condenser temperature (outdoor air temperature), the chiller should experience an approximate 2.5% increase in efficiency. In addition, their performance is poor compared to the performance of typical water-cooled chillers. This area is certainly an area of opportunity for energy savings in future construction, and is discussed further in the report. (The Appendix contains additional graphs that show monitored chiller efficiency.)

The air handler airflow delivery efficiencies were measured at 1367, 1375 and 1387 CFM/kW, which are below the design efficiencies by 40-60%. This is likely caused by increased pressure drop in the existing ductwork, which results in a decrease in airflow, compared to the standard testing conditions that are employed when fans are tested. Low pressure-drop duct design is important for achieving high air movement efficiencies.

VI. Energy Use – Data Center 6.2

DATA CENTER 6.2: ELECTRICAL EQUIPMENT AND BACKUP POWER SYSTEM

The facility utilizes an International Power Machine 160kVA uninterruptible power supply (UPS1), and a Chloride 50 Power Electronics 50kVA uninterruptible power supply (UPS2) to provide a constant supply of power of constant delivery voltage (480 V) to the data center. The UPS converts AC current and stores it as DC current in multiple battery packs. When the voltage is needed, it is converted back to AC current. In the event of a power loss, a 750 kW diesel generator will provide power for approximately 10 hours. Here as well, spot power measurements were taken at the UPS, both at the input and output in order to determine computer plug loads, as well as losses at the UPS system.

TABLE 6. UPS ELECTRICAL MEASUREMENTS

	Electrical Use ¹	Units
UPS1 Input	103.6	kW
UPS1 Output	96.3	kW
UPS1 Losses	7.3	kW
UPS1 Efficiency	93%	%
UPS2 Input	25.4	kW
UPS2 Output	22.8	kW
UPS2 Losses	2.6	kW
UPS2 Efficiency	90%	%

¹ Average measurement taken on 8/27/02, and 8/28/02.

Note, the UPS efficiencies at Data Center 6.2 are slightly higher than the efficiency measured for the UPS serving Data Center 6.1.

DATA CENTER 6.2: COOLING SYSTEM

The data center is cooled by a chilled water system that serves the data center, as well as several small computer labs. The chilled water system consists of two 220 Ton Trane rotary air-cooled chillers. The nominal efficiencies of the chillers are 1.3 kW/Ton.¹⁶ The chillers are piped in parallel, and both are typically operating at all times. The emergency distribution panel (EDP) serves one of the chillers. The chilled water pumps are 8.5 hp (hydraulic horsepower) pumps, and are constant speed. One main pipe feeds the cooling loads on each floor, however, the data center is the last load fed by the main pipe.

¹⁶ Based on 420 gpm, entering and leaving chilled water temperatures of 56 °F, and 44 °F, respectively, and entering condenser water temperature of 95 °F.

As with data center 6.1 , power consumption, flow, and chilled water temperatures¹⁷ were measured at each chiller over a period of several days. This was to determine the chiller efficiency over a period of varying temperatures.

Unlike the other data center, the chilled water feeds fan coil units (FCUs) in the ceiling plenum, which supplies the overhead duct system. The fan coil units are constant speed and have three –way valves. The system consists of a total of seven fan coil units, with cooling capacities ranging from 104,000 Btu/hr to 190,000 Btu/hr, and design airflow ranging from 5,300 cfm to 9,600 cfm. Air is returned through grills in the ceiling. Minimum outdoor air is brought in through the house air conditioning system. As with Data Center 6.1, there is no humidity control in Data Center 6.2.

The total chilled water load to all the FCUs was monitored using the technique of measuring flow rate, and pipe surface temperatures.¹⁸ As with the previous data center, it was necessary to identify the load solely to the data center, in order to segregate the chiller power consumption due to cooling of the data center only. The number and arrangement of the fan coil units did not allow for measurement of individual fan coil cooling load, nor air supply flow rate.

The spot measurements and average of trended measurements are listed in the table below. Please refer to the Appendix for graphs of the measurements over the entire monitored period. The chiller pump and chiller power are proportioned to the data center cooling load in order to properly determine the electrical end use in the data center.

¹⁷ These were measured using an Elite power-measuring instrument, an ultrasonic flow meter for pipe flow, and thermistors inserted in the Pete's plugs at the inlet and outlet of the chilled water line.

¹⁸ These measurements were made at the main branch that feeds only these units. Measurements of chilled water temperatures were performed by inserting thermistor probes between insulation and the pipe surface. Flow measurements were made using an ultrasonic flow meter.

TABLE 7. COOLING EQUIPMENT ELECTRICAL AND LOAD MEASUREMENTS

Equipment	Spot / Monitored	Date	Measurement	Units
Chiller Pumps - Total	Spot	9/4/02	23.5	kW
Chiller Pumps - Proportioned by Data Center Load	Spot	9/4/02	4.0	kW
Fan Coils (On circuits 23, 25, 27)	Spot	9/4/02	5.6	kW
Fan Coils (On circuits 29, 31, 33)	Spot	9/4/02	2.5	kW
Fan Coils (On circuits 35, 37, 39)	Spot	9/4/02	11.8	kW
DC Cooling Load From Chilled Water - From Monitoring of Chilled Water Use	Monitored	8/27/02 - 9/4/02	158.0	kW
DC Chiller kW From Monitoring - Average	Monitored	8/27/02 - 9/4/02	45.9	kW

DATA CENTER 6.2: LIGHTING

Lighting in the data center consists of T-8 tubular fluorescent lamps, and all lights were on when taking power measurements. Lighting Power: 2.65 kW (measured on 8/27/02) or 1.1 W/sf. These values are more than double what was measured for Data Center1.


DATA CENTER 6.2: SUMMARY MEASUREMENTS AND METRICS

The table below brings together all the equipment electrical measurements for the data center.

TABLE 8. SUMMARY OF ELECTRICAL MEASUREMENTS

Computer Loads	119.1	KW	59%
UPS Losses	9.9	KW	5%
HVAC - Air Movement	19.9	KW	10%
HVAC - Pumps and Chiller	49.0	KW	25%
Lighting	2.7	KW	1%
Total Energy Use	201.1	KW	100%

These results are shown graphically in the pie chart below.



The computer loads, based on the measured UPS power supply amounts to 59% of the data center energy usage. Pumping and cooling energy is the second largest consumer at 25%, and air movement from the fan coil units is 10%. Together, the HVAC component amounts to a significant 35% of data center energy use. Therefore, the HVAC components provide a significant opportunity for energy savings. Losses at the UPS consume 5% of the data center energy consumption. The percentage of lighting power consumption was the same for this data center, measured at 1%, though the energy density (W/sf) was higher.

Commensurate with the discussion under Data Center 6.1, different metrics are calculated for the data center energy use, and energy efficiency. To briefly reiterate, the computer load density is based on both gross area, which we equate to “data center floor area”, and on rack floor area. Both are extrapolated to 100% occupancy to predict future loads.

TABLE 9. ELECTRICAL CONSUMPTION METRICS

Data Center Gross Area	2,500	sf
Rack Area	432	sf (Calculated from a total of 83 racks, and area of 1 rack)
"Occupied" %	50%	Estimated from visual inspection.
Based on Gross Area:		
Computer Load Density	48	W/sf
Non-Computer Load Energy Density	33	W/sf
Projected Computer Load Density	95	W/sf
Based on Rack Area:		
Computer Load Density	276	W/sf
Projected Computer Load Density	551	W/sf
On an Individual Rack Basis:		
Computer Load Density	1.4	kW/Rack
Projected Computer Load Density	2.9	kW/Rack

The computer load density based on the data center area (gross area) is 48 W/sf. At full occupancy, the computer load density is projected to be 95 W/sf. This would require approximately 40 more tons of cooling, which based on the average measured chiller load, could be met by the chillers. The computer load density based on *rack area* is presently 276 W/sf, and is projected to be 551 W/sf at full occupancy. The average computer load, based on the number of racks is currently 1.4 kW/Rack, projected to be 2.9 kW/Rack at full capacity. The non-computer energy density, which includes HVAC, lighting, and UPS losses, is measured at 33 W/sf.

Commensurate with Data Center 6.1, the energy efficiency metrics are shown in the table below.

TABLE 10. HVAC EFFICIENCY METRICS

Metric	Value	Units
Cooling kW: Computer Load kW	0.58	--
Theoretical Cooling Load	40	Tons
Cooling Provided by Chilled Water and Fan Coil Units	44	Tons
Chiller 1 Efficiency	1.0	kW/ton
Chiller 2 Efficiency	1.1	kW/ton
Chiller 1,2 Design Efficiency ¹⁹	1.3	kW/ton
Average Chiller Efficiency	1.0	kW/ton
Fan Coil Unit Design Efficiency	2,370	CFM/kW
Overall HVAC Efficiency	1.6	kW/ton

From the above table it is shown that the “cooling efficiency” of approximately 0.6 kW/kW is significantly less efficient than the cooling efficiency for Data Center 6.1. This can be explained by the differences in equipment, but is not an entirely valid comparison, since Data Center 6.1’s metrics suggests that significant cooling was provided by the whole building air conditioning system. This does not appear to be the case with Data Center 6.2, where the measured cooling load was more than 10 tons larger than the theoretical cooling load.²⁰

The performance of the chillers is similar to what was observed with Data Center 6.2's chillers. (i.e., The performance was slightly better than ARI rated performance, which is expected for the operating conditions.) However, the performance of water-cooled chillers far outweighs the performance of these units, and is an opportunity for energy savings in future construction.

The design efficiencies of the FCUs are comparable to the design efficiencies of the AHUs used in Data Center 6.1, though the actual efficiencies were not measured.

¹⁹ The nominal efficiencies cannot be directly compared to the average operating efficiencies, since the nominal efficiencies are based on full load capacities, and the specific conditions cited previously.

²⁰ This can be attributed to measurement error of the cooling load, and that computer loads were assumed to be constant, while they may vary a small percent over time. This assumes, per the drawings, no other fan coil units on the first floor serve non data center rooms, which if present, would explain the small difference.


VII. Energy Efficiency Recommendations

GENERAL GUIDELINES FOR FUTURE CONSTRUCTION

Efficient Chilled Water System

Water-cooled chillers offer enormous energy savings over air-cooled chillers, particularly in dry climates, such as the bay area because they take advantage of evaporative cooling. Since lower temperature media is cooling the chiller, it can reject heat more easily, and does not have to work as hard. Though the addition of a cooling tower adds maintenance costs associated with the water treatment, we have found that the energy savings outweigh the maintenance costs. Within the options of water cooled chillers, variable speed centrifugal are the most energy efficient, because they can operate very efficiently at low loads. The graph below compares the energy performance of various chiller types.

Comparison of Typical Chiller Efficiencies over Load Range

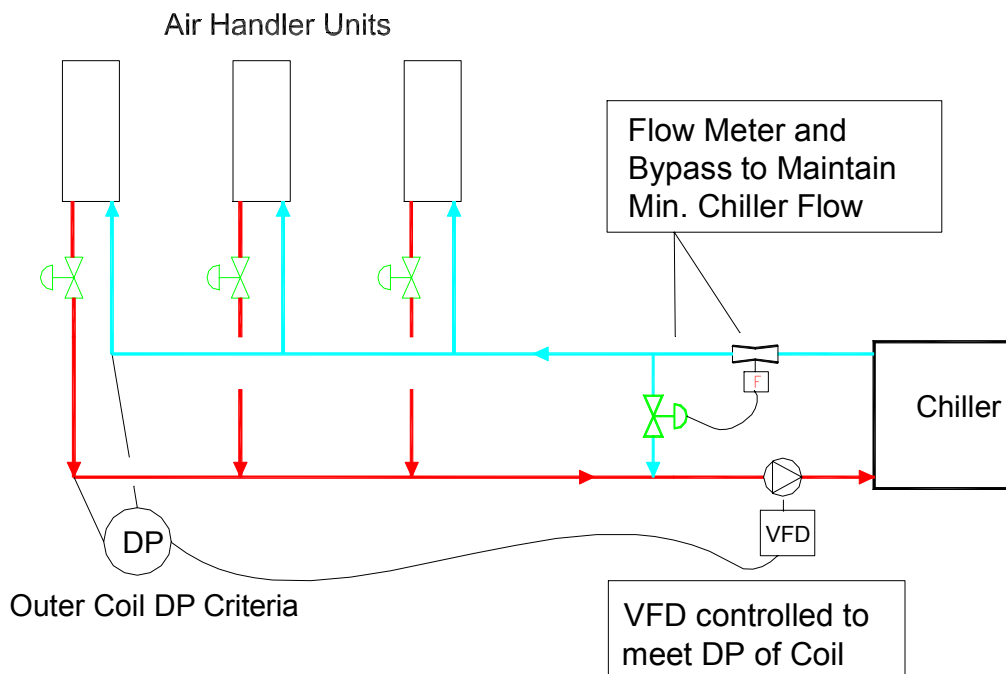


Chiller 1	250-Ton, Screw, Standard Efficiency, Air Cooled
Chiller 2	216-Ton, Screw, Water Cooled
Chiller 3	227-Ton, Centrifugal, Constant Speed, Water Cooled
Chiller 4	227-Ton, Centrifugal, Variable Speed, Water Cooled

Though there are efficient air cooled chillers, the larger size of water cooled chillers has resulted in more care given to efficiency and life cycle costs compared to air cooled chillers.

The selection of the auxiliary equipment, including the cooling tower, pumps, and pumping strategy should also be considered carefully. For example, variable speed fans on cooling towers allow for optimized cooling tower control. Premium efficiency motors and high efficiency pumps are recommended, and variable speed pumping is a ripe opportunity for pump savings. Variable pumping strategies can be achieved in a primary/secondary scheme, where the primary pumps operate at constant speed and directly feed water to the chiller, and the secondary pumps are variable speed and serve the air-handling units. A more energy efficient scheme is primary-only variable speed pumping strategy. Pumping savings are based on the cube law: the cube of the reduction in pump speed reduces pump power, which is directly proportional to the amount of fluid pumped.

A primary only variable pumping strategy must include a bypass valve that ensures minimum flow to the chiller, and the use of two-way valves at the air-handling units in order to achieve lower pumping speeds. The control speed of the bypass valve should also meet the chiller manufacturers recommendations of allowable turndown, such that optimum chiller efficiency is achieved.²¹ The diagram below describes the primary-only variable speed pumping strategy.



²¹ This basically means that the flow through the chiller should be varied slow enough such that the chiller is able to reach a quasi-steady state condition and able to perform to its maximum efficiency.

Air Management

The standard practice of cooling data centers employs an underfloor system fed by CRAC units. There are a number of potential problems with such systems: an underfloor system works on the basis of thermal stratification. This means that as the cool air is fed from the underfloor, it absorbs energy from the space, warming up as a result, and rises. In order to take advantage of thermal stratification, the return air must be collected at the ceiling level. CRAC units often have low return air grills. Though there are CRAC units available with return grills located on top, the unit may not be tall enough to take advantage of thermal stratification. As a result, the CRAC units are often re-circulating cool or only moderately warmed air. Furthermore, they are often located along the perimeter of the building, and not dispersed throughout the floor area, where they can more effectively treat warm air. One alternative is to install transfer grills from the ceiling to the return grill. Another common problem with underfloor supply is that the underfloor becomes congested with cabling, increasing the resistance to air flow. This results in an increase in fan energy use. A generous underfloor depth is essential for effective air distribution (we have seen 3 feet in one facility).

An alternative to underfloor air distribution is high velocity overhead supply, combined with ceiling height return. A central air handling system can be a very efficient air distribution unit. Design considerations include using VFDs on the fans, low-pressure drop filters, and coils. An additional advantage of a central air handling system is that it can be specified with an economizer function. With the favorable climate in the Bay Area, economizing can reduce the cooling load for a majority of the hours of the year.

Another common problem identified with CRAC units is that they are often fighting each other in order to maintain a constant humidity set point. Not only is a constant humidity set point unnecessary for preventing static electricity (the lower limit is more important), but also it uses extra energy. A central air-handling unit has a better ability to control overall humidity than distributed CRAC units.

Air Management – Rack Configuration

Another factor that influences cooling in data centers is the server rack configuration. It is more logical for the aisles to be arranged such that servers' backs are facing each other, and servers' fronts are facing each other. This way, cool air is drawn in through the front, and hot air blown out the back. The Uptime Institute has published documents describing this method for air management.²² Our observations of both data centers showed an inconsistent rack configuration.

Commissioning of New Systems and Optimized Control Strategies

Many times the predicted energy savings of new and retrofit projects are not fully realized. Often, this is due to poor and/or incomplete implementation of the energy efficiency recommendations. Commissioning is the process of ensuring that the building systems perform as they were intended to by the design. Effective commissioning

²² <http://www.upsite.com/TUIpages/whitepapers/tuiaisles.html>

actually begins at the design stage, such that the design strategy is critically reviewed. Either the design engineer can serve as the commissioning agent, or a third party commissioning agent can be hired. Commissioning differentiates from standard start-up testing in that it ensures systems function well relative to each other. In other words, it employs a systems approach.

Many of the problems identified in building systems are often associated with controls. A good controls scheme begins at the design level. In our experience, an effective controls design includes 1) a detailed points list, with accuracy levels, and sensor types, and 2) a detailed sequence of operations. Both of these components are essential for successfully implementing the recommended high efficiency chilled water system described above.

Though commissioning is relatively new to the industry, various organizations have developed standards and guidelines. Such guidelines are available through organizations like the Portland Energy Conservation Inc., at www.peci.org, or ASHRAE, Guideline 1-1996.

Lighting Controls

The lighting power and lighting power densities for Data Center 6.2 were more than twice those of Data Center 6.1. This is likely from occupants/engineers entering the Data Center, and turning the lights on. Lighting controls, such as occupancy sensors may be appropriate for these types of areas that are infrequently, or irregularly occupied. If 24 hour lighting is desired for security reasons, scarce lighting can be provided at all hours, with additional lighting for occupied periods.

DATA CENTER 6.1 SPECIFIC OBSERVATIONS

Verification of Bypass Control: The chilled water pumping for Data Center 6.1 utilizes a primary only, variable speed drive (VSD) system, with a bypass control valve. From our observation of the EMCS, the VSD is being controlled via a differential pressure (dP) sensor, however the control scheme for the bypass valve is not clear. A pressure-independent bypass control is the most effective, where the actual flow supplied to the chiller is monitored, and used as the control input to the bypass control valve. A pressure-dependent system will maintain a constant differential pressure, and is controlling flow by using pressure as a surrogate. We suggest that the control scheme for the bypass control valve be examined to ensure that it is being controlled properly.

Three - Way Valves and Bypass: Though primary-only, variable pumping system equipment has been installed, it is not clear whether the air handling units serving the data center and fan coil units serving the computer labs are equipped with two-way valves, as they should be. In order for a variable system to function as intended, the air handling units and fan coil units should be equipped with two way control valves.

Chiller Staging: Constant speed chillers are designed to operate more efficiently at their nominal loads. Currently, both chillers are running most of the time, regardless of the load. (See graphs in Appendix.) It would be more efficient to stage the chillers such that the smaller chiller comes on when the larger chiller is unable to satisfy the cooling

requirements. This staging could be based on the primary chiller being unable to meet its chilled water set point. The measured data showed that the load did not exceed 90 tons, and therefore the large chiller should be capable meeting the load most of the time. Attention should be given to how quickly flow is diverted from the primary chiller so that it does not go off inadvertently on low load.

Triple Duty Valves: Triple duty valves have been installed on the discharge of each of the chilled water pumps. We recommend that the triple duty valves be opened completely.

DATA CENTER 6.2 SPECIFIC OBSERVATIONS

Chiller Oscillations: The measured data identified power oscillations with chiller 1. This could be due to cycling of one of the compressors. The controls of this chiller should be investigated, since this cycling effect has an adverse effect on energy consumption and will increase maintenance cost. Though chiller staging is achievable for Data Center 6.1, the measured data shows that the chilled water load for the building hosting Data Center 2 exceeds the nominal load of one chiller.

Close 4 Inch Bypass: The mechanical drawings show the existence of a 4-inch bypass on the chilled water loop, located on the first floor. Visual observation of the fan coil units shows the existence of three-way valves (though this differs from the mechanical drawings). Upon confirmation of three-way valves on all fan coil units, this bypass can be closed.


Primary- Only Variable Speed Conversion: The current constant speed pumping strategy could be converted to a variable speed system by installing VSDs on the pumps, installing a controlled bypass line to ensure minimum flow through the chillers, and by converting the three-way valves to two-way valves. Note, this is the system that is already installed on the chilled water system serving Data Center 6.1.

High Velocity Diffusers and Air Management: Both data centers utilize overhead air supply. Diffusers should be sized for high velocities such that air is directed downwards in aisles facing the fronts of the servers. Also see Air Management - Rack Configuration.

Triple Duty Valves: Triple duty valves have been installed on the discharge of each of the chilled water pumps. We recommend that the triple duty valves be opened, and that the pump impellers be trimmed for balancing. This has the same effect as reducing the pump size and flow, without sacrificing efficiency. If a conversion is made to variable speed pumping, then the impeller does not have to be trimmed.


APPENDICES – MONITORED DATA – FACILITY 6, DATA CENTER 6.1

Facility 6 Data Center 6.1 Total Chiller Characteristics




Facility 6 Data Center 6.1


Chiller Flow




Facility 6 Data Center 6.1
100 Ton Chiller - Chilled Water Temperatures




Facility 6 Data Center 6.1 100 Ton Chiller Characteristics




Facility 6 Data Center 6.1
40 Ton Chiller - Chilled Water Temperatures



Facility 6 Data Center 6.1 40 Ton Chiller Characteristics




Facility 6 Data Center 6.1 Data Center Chilled Water Flow




Facility 6 Data Center 6.1

Data Center Chilled Water Temperatures




Facility 6 Data Center 6.1
Outside Dry Bulb Air Temperature




APPENDICES – MONITORED DATA – FACILITY 6, DATA CENTER 6.2


Facility 6 Data Center 6.2 Total Chiller Characteristics




Facility 6 Data Center 6.2 Combined Chiller Flow




Facility 6 Data Center 6.2
Chiller 1 - Chilled Water Temperatures




Facility 6 Data Center 6.2 Chiller 1 Characteristics




Facility 6 Data Center 6.2 Chiller 1 Characteristics




Facility 6 Data Center 6.2
Chiller 2 - Chilled Water Temperatures




Facility 6 Data Center 6.2 Chiller 2 Characteristics




Facility 6 Data Center 6.2
Data Center CHW Flow Rate




Facility 6 Data Center 6.2
Data Center - Chilled Water Temperatures and Load



Facility 6 Data Center 6.2
Data Center Chilled Water Temperatures and Load



Facility 6 Data Center 6.2
Outside Air Temperature



**Appendix VIII. “Data Center Energy Benchmarking Case Study,
Facility 7”; PIER Project, Rumsey Engineers, Inc.; LBNL –
William Tschudi, 2003**

DATA CENTER ENERGY BENCHMARKING CASE STUDY

FEBRUARY 2003



FACILITY 7
SPONSORED BY:



LAWRENCE BERKELEY NATIONAL LABORATORY
PREPARED BY:



99 LINDEN STREET
OAKLAND, CA 94607
(510) 663-2070

Acknowledgements

Rumsey Engineers is grateful to the facility managers/directors, engineers and operators for their generous assistance and cooperation. Special thanks to Christine Condon of PG&E for providing monitoring equipment on short notice. Thanks also to the Lawrence Berkeley National Laboratory, and the California Energy Commission for funding this project.

Disclaimer

Neither Rumsey Engineers, LBNL nor any of its employees makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any data, information, method, product or process disclosed in this document, or represents that its use will not infringe any privately-owned rights, including but not, limited to, patents, trademarks, or copyrights.

Contents

I. EXECUTIVE SUMMARY	1
II. DEFINITIONS.....	7
III. INTRODUCTION.....	9
IV. SITE OVERVIEW.....	10
V. ENERGY USE.....	11
ELECTRICAL EQUIPMENT AND BACKUP POWER SYSTEM	11
COOLING SYSTEM	12
LIGHTING	17
DATA CENTER ELECTRICITY END USE.....	17
WHOLE BUILDING ELECTRICITY END USE	19
HVAC EFFICIENCY METRICS	23
VI. ENERGY EFFICIENCY RECOMMENDATIONS.....	26
DEDICATED CHILLER FOR DATA CENTER AND OTHER CONTROL MEASURES	26
CONVERSION TO VARIABLE SPEED PUMPING ON SECONDARY LOOPS AND OTHER	
PUMPING SAVINGS.....	28
COOLING TOWER STAGING	30
ECONOMIZER BASED COOLING	30
AIR MANAGEMENT	31
UPS REPLACEMENT.....	32
COMMISSIONING OF NEW SYSTEMS AND OPTIMIZED CONTROL STRATEGIES.....	32
OPTIMIZATION OF EMCS.....	33
LIGHTING CONTROLS	33
IN HOUSE ENGINEERING RECOMMENDATIONS.....	26

APPENDICES

A -- CHARTS OF MONITORED DATA

I. Executive Summary

Rumsey Engineers and the Lawrence Berkeley National Laboratory (LBNL) have teamed up to conduct an energy study as part of LBNL's Data Center Load Characterization and Roadmap Project, under sponsorship by the California Energy Commission (CEC). This study will aid designers to make better decisions about the design and construction of data centers in the near future. Data centers at four different organizations in Northern California were analyzed during the period of September 2002 to December 2002, with the particular aim of determining the end-use of electricity.

This report documents the findings for one of the case studies – termed Facility 7. Additional case studies and benchmark results as they become available will be provided on LBNL's website (<http://datacenters.lbl.gov>). For comparison purposes, the results of a similar benchmarking study completed for the Pacific Gas and Electric Company (PG&E) in 2001 are included in this report.

Facility 7 contains two data centers, in two separate floors, in a large office building. The facility is a financial institution, and has a variety of data equipment, that includes file servers, tape storage robots, and printers. An addition, another floor contains check-processing equipment, which is also served by the critical facility equipment, but for purposes of this study, was excluded where possible. Only a portion of the data center resembles the server farms that became common as a result of the Internet Age.¹ The data center gross area is approximately 74,000 square feet (sf), while the entire building is 1.4 million sf. The data center electricity and the building electricity end use are evaluated.

The whole building and data center are served by a chilled water plant. Primary chilled water directly feeds all the building's main air handlers. A heat exchanger separates the primary chilled water loop from the secondary chilled water loop, which supplies chilled water to the computer room air conditioners (CRAC) units. The CRAC units pressurize a raised floor, and the air handlers supply air through a VAV system overhead.


The current computer energy loads are listed in the table below. A qualitative estimate of the loading of the racks was made, and the future computer energy loads were estimated based on this loading. For comparison purposes the computer loads of another data center studied in this project (CEC funded) and other data centers studied in the PG&E project are also included. The computer loads are also shown graphically.

¹ Based on the rack configuration, high density of computers, and absence of the large mainframe servers that were common in older data centers.

CURRENT AND FUTURE COMPUTER LOADS

Data Center	Data Center Area (sf)	Computer Load (kW)	Computer Load Energy Density (W/sf)	Occupancy (%)	Projected Computer Load Energy Density (W/sf)
Data Center 7	74,000	1,395	19	80%	24
Data Center 1	62,870	1,500	24	75%	32
Data Center 2	60,400	2,040	34	65%	52
Data Center 3	25,000	1,110	44	85%	52
Data Center 6.1	2,400	155	65	80%	81
Data Center 6.2	2,501	119	48	50%	95
Data Center 8.1	26,200	222	8	30%	27
Data Center 8.2	73,000	1,059	15	30%	50


Computer Energy Loads



The measured computer load densities at Facility 7 are significantly smaller than the computer load densities measured in the previous study. The measurement projects a full occupancy density of 24 W/sf, which is well below all of the full occupancy densities projected for the other data centers.


The remaining energy loads of the data center include chiller, and chilled water plant energy (proportioned to the data center load), CRAC unit power, lighting, and uninterruptible power supply inefficiencies. Due to the critical nature of the facility, an efficiency of the UPSs were not obtained, but an efficiency was assumed, based on values observed at another site.

The data center electrical end use is shown below in graphical format, and is listed in tabular format in the report.




A large percentage, approximately 47%, of the total electrical load is from the computer loads. However, the HVAC loads contribute a significant percentage at 37%. Therefore, efficiency improvements could result in significant energy savings. In addition, the estimated lighting and UPS consumption represent an opportunity for energy savings, where redundancy requirements permit such changes in operation. These are discussed in detail in the report.

The whole building electricity end use was also determined, and is shown in two formats. The first, separates the data center loads from the non-data center loads, and the second is categorized based on major equipment categories. This data is included in tabular format in the report.



The whole building consumes an average of 5 MW of electricity. The major consumers are 1) the data center computer loads, 2) the office plug loads, lighting loads, and miscellaneous loads (which include elevator loads), and 3) the data center HVAC. They are approximately 30%, 30%, and 20%, respectively, or 70% together. The data center alone contributes to 62% of the total building energy.



The largest consumer is the UPS power at approximately 1.9 MW, or 38% for weekend operation, and 33% for weekday operation. The difference between this value, and the computer loads cited in the previous paragraph is due to the UPS losses. The total HVAC power for the whole building is approximately 30% for both scenarios, and the office, lighting, elevators and miscellaneous loads account for 34%. This representation further emphasizes the electrical consumption of HVAC equipment, and the relevance of energy efficiency measures. As shown earlier, approximately 2/3 of the HVAC power can be attributed to the data center cooling exclusively. The category of lighting, office, elevators, and miscellaneous sources emphasizes that lighting, as well as power management within offices is important, though this is not a focus of the study.

The performance of the HVAC system can be evaluated based on energy efficiency metrics. Though the cooling power is represented in W/sf, a more useful metric for evaluating how efficiently the data center is cooled can be represented as a ratio of cooling power to computer power. This essentially removes the variable of how tightly packed the computers are. The more traditional metrics of energy per ton of cooling (kW/Ton) are calculated for individual chillers, total chilled water plant (chillers, cooling towers, pumps) and the data center cooling efficiency. The data center cooling efficiency includes the chilled water plant power weighted by the data center load, data center air handler and CRAC unit power.

FACILITY 7 EFFICIENCY METRICS

Metric	Value	Units
Building Energy Density	3.6	W/sf
Data Center Computer Power Density	18	W/sf

Metric	Value	Units
Data Center Cooling Power Density	15	W/sf
Non Computer Load Density	21	W/sf
Cooling kW / Computer Load kW	0.7	--
Chiller 3 Efficiency	0.9	kW/ton
Chiller 4 Efficiency	0.6	kW/ton
Chilled Water Plant Efficiency	1.1	kW/ton
Total Data Center HVAC Efficiency	1.7	kW/ton
Theoretical Cooling Load *	553	Tons
Cooling Provided by CRAC Units	520	Tons
Cooling Provided by Air Handlers	121	Tons
Measured Cooling Load	641	Tons

* Based on computer loads, lighting loads, and fan energy.

The data center computer load density is small, relative to what is observed at other data centers. Hence, the cooling energy density is also small, at 15 W/sf, relative to other facilities. The “cooling efficiency”, which is the efficiency normalized to the computer power is 0.7 Cooling kW/ Computer kW. This means that for 1 kW of energy input, only 1.4 kW of energy is removed. This value is slightly higher than the measured efficiencies of 0.5 kW/kW and 0.6 kW/kW at two other monitored data centers, which utilizes air cooled chillers and fan coil units, and air cooled CRAC units, respectively. Another monitored site has an efficiency of 1.3 kW/kW, which utilizes a water cooled reciprocating chiller and computer room air handlers with humidification and reheat. Though a water cooled chiller plant could operate extremely efficiently, it will not if the fundamental equipment, air delivery and pumping systems are inefficient.

Chiller efficiencies were obtained for chillers 3 and 4, and on average are 0.9 and 0.6 kW/ton, respectively. Both are water cooled, centrifugal constant speed chillers. The efficiency of chiller 4 meets the rated efficiency of 0.6 kW/ton, but the observed efficiency should be better, since the operating conditions are more favourable (at 68 °F entering condenser water temperature).²

Several opportunities for energy savings, addressing the chiller plant, and other areas are described in detail in the “Energy Efficiency Recommendations” section of the report.

² The rated conditions are: 80 °F entering condenser water temperature, chilled water setpoint of 42 °F. The operating conditions are: 65 °F entering condenser water temperature, chilled water setpoint of 42 °F.

II. Definitions

Data Center Facility	A facility that contains both central communications equipment, and data storage and processing equipment (servers) associated with a concentration of data cables. Can be used interchangeably with Server Farm Facility
Server Farm Facility	A facility that contains both central communications equipment, and data storage and processing equipment associated with a concentration of data cables. Can be used interchangeably with Data Center Facility. Also defined as a common physical space on the Data Center Floor where server equipment is located (i.e. server farm)
Data Center Floor / Space	Total footprint area of controlled access space devoted to company/customer equipment. Includes aisleways, caged space, cooling units, electrical panels, fire suppression equipment, and other support equipment. Per the Uptime Institute Definitions, this gross floor space is what is typically used by facility engineers in calculating a computer load density (W/sf). ³
Data Center Occupancy	This is based on a qualitative estimate on how physically loaded the data centers are. For calculations, the facility's assessment of this value is utilized.
Data Center Cooling	Electrical power devoted to cooling equipment for the Data Center Floor space
Data Center Server/Computer Load	Electrical power devoted to equipment on the Data Center Floor. Typically the power measured upstream of power distribution units or panels. Includes servers, switches, routers, storage equipment, monitors, and other equipment.

³ Users look at watts per square foot in a different way. With an entire room full of communication and computer equipment, they are not so much concerned with the power density associated with a specific footprint or floor tile, but with larger areas and perhaps even the entire room. Facilities engineers typically take the actual UPS power output consumed by computer hardware and communication equipment in the room being studied (but not including air handlers, lights, etc.) and divide it by the gross floor space in the room. The gross space of a room will typically include a lot of areas not consuming UPS power such as access aisles, white areas where no computer equipment is installed yet, and space for site infrastructure equipment like Power Distribution Units (PDU) and air handlers. The resulting gross watts per square foot (watt/ft²-gross) or gross watts per square meter (watt/m²-gross) will be significantly lower than the watts per footprint measured by a hardware manufacturer in a laboratory setting.

Computer/Server Load Measured Energy Density	Ratio of actual measured Data Center Server Load in Watts (W) to the square foot area (ft ² or sf) of Data Center Floor. Includes vacant space in floor area
Computer Load Density – Rack Footprint	Measured Data Center Server Load in Watts (W) divided by the total area that the racks occupy, or the rack “footprint”.
Computer Load Density per Rack	Ratio of actual measured Data Center Server Load in Watts (W) per rack. This is the average density per rack.
Computer /Server Load Projected Energy Density	Ratio of forecasted Data Center Server Load in Watts (W) to square foot area (ft ² or sf) of the Data Center Floor if the Data Center Floor were fully occupied. The Data Center Server Load is inflated by the percentage of currently occupied space.
Cooling Load Tons	A unit used to measure the amount of cooling being done. Equivalent to 12,000 British Thermal Units (BTU) per hour.
Chiller Efficiency	The power used (kW), per ton of cooling produced by the chiller.
Air Handler Efficiency 1	The air flow (CFM) per power used (kW) by the CRAC unit fan
Air Handler Efficiency 2	The power used (kW), per ton of cooling achieved (ton) by the air handling unit.
Cooling Load Density	The amount of cooling (tons) in a given area (ft ² or sf)
Air Flow Density	The air flow (CFM) in a given area (ft ² or sf)

III. Introduction

This report describes the measurement methodology and results obtained for this case study. The facility is a large financial facility that includes office space and data centers. The data centers were measured collectively. Electricity end use for the entire building and data centers is determined. This was achieved through a combination of spot electrical measurements, temperature and flow measurements on various mechanical equipment, spot measurements of utility meters for computer loads, miscellaneous office loads, and trended data on mechanical systems from the Energy Management Control System (EMCS). The computer load density is also determined based on the gross area of the data center, as this number, in watts per square foot (W/sf) is the metric typically used by facility engineers to represent the computer power density. Based on the owner's assessment of the data center occupancy, the computer load density at full occupancy is extrapolated.

Additional information was collected, where necessary, in order to determine the operating efficiencies of the cooling equipment. These efficiencies are compared to the design efficiencies. Opportunities for energy efficiency improvements are described, which are based on observation of the mechanical system design, and measured performance.

IV. Site Overview

Facility 7 is a large financial institution located in San Francisco, California. The building has a gross area of 1.4 million square feet (sf). It consists of several floors of office area, and three floors that are dedicated to computer equipment, and check processing. The data center gross floor area is approximately 74,000 sf, and consists of approximately 5% of the total building area.⁴ The data centers host a variety computer equipment that includes servers and networking equipment, mainframe computers, tape storage robots, and printers. A portion of the data centers is arranged in the typical server farm rack style.



Inside Data Center

The data centers are operated 24 hours a day. The check processing areas are not considered to be typical data center equipment, and are therefore not included in the majority of the calculations. In order to avoid including this data, weekend measurements were obtained as much as possible during the measurement period.

The data centers are cooled both by an underfloor system and overhead system. The underfloor system is supplied cool air by water-cooled computer room air conditioners (CRACs)⁵ and the overhead system by typical office air handlers. The central cooling plant serves the entire building.

⁴ Note, this is not the total floor area of the floors housing the computer rooms, only the computer room area.

⁵ Termed “environmental air conditioners” or EACs by the facility.

V. Energy Use

ELECTRICAL EQUIPMENT AND BACKUP POWER SYSTEM

The facility utilizes an Exide 2225 kVA uninterruptible power supply (UPS) and a Teledyne 2500 kVA UPS. The UPSs provide a constant supply of power to the data center at constant delivery voltage (480/277 V). The UPS converts AC current and stores it as DC current in multiple battery packs. When the voltage is needed, it is converted back to AC current. In the event of a power loss, four 3 megawatt (MW) Diesel generators can provide power for approximately 8 days at maximum generator load.

Spot power measurements were taken at the UPSs by reading the instantaneous power draw at the utility meters. In order to avoid including the power draw by the check processing equipment, several readings were taken during weekend operation, after confirming with facility personnel that check processing would not be active. The output of the UPSs were not measured, as this would involve an electrical hook up to critical facility equipment. For the purposes of estimating UPS heat losses, an efficiency of 80% was assumed. This was based on experience at other facilities where power measurements were conducted by measuring at the input and output sides of the UPS.⁶



Teledyne UPS

The most commonly used metric among mission critical facilities is the computer load density in watts consumed per square foot (W/sf). However, the square footage is not always consistent between designers. This inconsistency has been a problem.⁷ Some data centers use kVA/rack or kW/rack as a design parameter. Our definition of “Data Center Floor Area” includes the gross area of the data center, which includes rack spaces, aisle spaces, and areas that may eventually contain computer equipment. Per the Uptime Institute, the resulting computer load density (W/sf) is consistent with what facility

⁶ Measurements at other facilities indicated an efficiency of 78% for an Emerson Accupower 500 kVA installed approximately 20 years ago. The measured UPS was approximately 30% loaded. The UPSs at Facility 7 are also under-loaded at 40%.

⁷ See “Data Center Power Requirements: Measurements from Silicon Valley”, by Mitchell-Jackson, Koomey, Nordman, & Blazek, December 2001. It is available on the web at http://enduse.lbl.gov/Info/Data_Center_Journal_Article2.pdf.

engineers use, though this is different from the “footprint” energy density that manufacturers use. The data center floor area was estimated from drawings by the in-house mechanical engineering company, and is 74,000 sf, and 104,000 sf when check processing areas are included.⁸ The UPS data, estimated UPS losses, and computer densities are listed in the table below.

TABLE 1. UPS ELECTRICAL MEASUREMENTS

	Sep 14 (Sat)	Sep 28 (Sat)	Sep 29 (Sun)	Oct 1 (Tue)	Oct 5 (Sat)	Oct 6 (Sun)
UPS 1 Input (Exide)	1056	861	862	1112	976	980
UPS 2 Input (Teledyne)	820	765	780	820	808	808
Total UPS Input (kW)	1876	1626	1642	1932	1784	1788
Calculated UPS Output (kW)	1501	1301	1314	1546	1427	1430
Calculated UPS Losses	375	325	328	386	357	358
Computer Density (W/sf)	20	18	18	15	19	19
Projected Computer Density (W/sf)	24					

The total UPS input power varies from 1626 kW to 1932 kW, with the peak power on a the only weekday measurement. The check processing rooms, which are variable loads are likely to account for this difference between weekday and weekend operation. However, the deviation from weekday to weekend operation is relatively small at 200 kW, or 10%. Note, the weekday measurements were taken during working hours in the early afternoon. This suggests that the check processing is additional load is an insignificant, or is highly variable even during weekday, working hours. The weekend measurements vary between 1626 kW and 1876 kW. The computer load density varies from 15 W/sf to 20 W/sf, with the minimum corresponding to the weekday operation. This value is small compared to measurements made at other facilities, which have computer load densities of 30 – 50 W/sf.

The Data Center Occupancy is a qualitative estimate of how physically full the rooms are. Per a meeting with the facility personnel, the approximate occupancy of each floor was obtained, and a geometrically weighted average occupancy was estimated. Based on this rough occupancy of 80%, the fully loaded computer load density, excluding the check processing areas, is projected on average to be 24 W/sf.⁹

COOLING SYSTEM

The facility has a central plant that serves both the data centers and office areas. It consists of four constant speed centrifugal chillers. Three of these were installed in 1973

⁸ The check processing areas are included for calculating the computer load density for the weekday measurement obtained on October 1.

⁹ Occupancy, and square footage data based on estimate given by facility.


and have a capacity of 1500 tons.¹⁰ During an expansion that occurred in 1994, a fourth chiller was added with a capacity of 1380 tons.¹¹ The purpose of the expansion was to split the chilled water plant into two side, for redundancy purposes.

The chillers are cooled by eight variable speed drive cooling towers, which are forced draft and are located indoors. These are typically operated in groups of four. The chilled water setpoint is 42 °F, and the condenser water setpoint is 68 °F. The chilled water is supplied in two directions by primary pumps. The main loop is fed to eight large air handlers, and the UPS room CRAC



Chiller 4

units. There are four primary pumps, three originally installed, and one of which was installed during the expansion. They are driven by 150 horsepower (hp) and 125 hp



Intercoolers


motors. One primary chilled water pump is typically on. There are five condenser water pumps, three older, and two newer, 75 and 60 hp. One condenser pump is typically on. The secondary loop that feeds the CRAC units is separated from the primary loop by two shell and tube heat exchangers.

¹⁰ Per mechanical schedule. Based on a evaporator flow rate of 3000 gpm, entering and leaving chilled water temperatures of 54 °F and 42 °F, respectively.

¹¹ Per mechanical schedule. Based on a evaporator flow rate of 2756 gpm, entering and leaving chilled water temperatures of 54 °F and 42 °F, respectively.

There are six secondary loop pumps, all 125 hp. A 200,000 gallon thermal storage tank provides backup chilled water if needed, but is typically not used. As mentioned, the chilled water plant is split into two sides that are mirror images. Each side includes two chillers, one shell and tube heat exchanger, two primary chilled water pumps, two/three condenser water pumps, and three secondary loop pumps.

The CRAC units, each have a 20 ton capacity. They are supplied chilled water by the secondary loop. A bypass valve controls flow to maintain a differential pressure in the secondary loop. Each CRAC has two way control valves. The units control to a return temperature setpoint of 70 °F, and relative humidity of 50% \pm 5%. Humidity control has been disabled on several of the units. There are two overhead air handlers that supply 55 °F air to the data center areas. Each has an air flow capacity of 170,000 cubic feet per minute (cfm), and cooling capacity of 6,227,000 British Thermal Units per hour (Btu/hr) or 520 tons. The overhead air handlers mix a minimal amount of outdoor air with return air, and do not have motorized dampers for outdoor air economizing. Steam is supplied for humidification.



Computer Room Air Conditioner

Spot power measurements were obtained on the pumps, air handler fans, and CRAC units using a power meter (PowerSight). Long term power monitoring was setup on the cooling tower fans, chillers, and CRAC units over a period of four weeks. Since the chiller load serves office air handlers, data center air handlers, and the CRAC units, it was necessary to identify the chilled water supplied solely to the data center, in order to segregate the chiller power consumption due to cooling of the data center only. Consequently, the chilled water supply, return, and flow were monitored to each air handler. When combined with the secondary chilled water load, the total chilled water load to the data

center can be determined. Where possible, data was obtained from the EMCS. This included individual chiller flow, chilled water supply and return temperatures, secondary loop chilled water flow, secondary loop chilled water supply and return temperatures, and chiller current. The monitored chiller current is for one phase, and is converted to power using the monitored power factor. Using this information, the chiller efficiency, total chilled water plant efficiency, and data center cooling efficiencies were determined.

The spot measurements, and average of trended and monitored points are listed in the table below. Please refer to the Appendix for graphs of the measurements over the entire monitored period, and tables with measurements for corresponding to each weekend day of operation during the measurement period. The "Data Center HVAC Pumps, Chiller, Cooling Tower" includes the chiller pump, chiller, and cooling tower power proportioned to the data center cooling load. The "Data Center HVAC Air Movement" power includes the total power for the data center dedicated overhead air handlers, as well as the CRAC units' power. The CRAC units' power proportioned to the data center load, versus non data center CRAC unit load in order to properly determine electrical end use for the data center.¹²

TABLE 2. COOLING EQUIPMENT ELECTRICAL AND LOAD MEASUREMENTS

Equipment	Spot / Monitored / Trended	Average Weekend Measure- ment	Weekday (Oct 1) Measure- ment	Units
HVAC Equipment Electrical Measurements				
Chiller 4	Monitored, Modified Trended Data	554.6	--	kW
Chiller 3	Monitored, Modified Trended Data	795.9	801.0	kW
Cooling Towers	Monitored	55.1	35.1	kW
Primary Chilled Water Pump (No. 1)	Spot	111.6	--	kW
Condenser Water Pump (No. 23)	Spot	47.8	--	kW
Secondary (Environment) Chilled Water Pump (No. 20)	Spot	91.7	--	kW
Air Handler 4	Spot and Monitored	41.3	--	kW
Air Handler 5	Spot	40.8	--	kW
CRAC Units (EACs) Total	Monitored	350.6	412.2	kW

¹² All CRAC units condition computer room equipment, but for the CRAC unit that conditions the Teledyne UPS room.

Equipment	Spot / Monitored / Trended	Average Weekend Measure- ment	Weekday (Oct 1) Measure- ment	Units
Cooling Load Measurements				
Chiller Tonnage	Trended	861	968	Tons
Cooling Provided by CRACs - Computer Room Cooling	Monitored	511	563	Tons
Cooling Provided by Air Handlers 4 and 5	Monitored	120	126	Tons
Data Center - Total Measured Cooling Load	Monitored	632	689	Tons
Chiller Efficiency				
Chiller 4 Efficiency	Monitored	0.6		kW/Ton
Chiller 3 Efficiency	Monitored	0.9		kW/Ton
Data Center Attributed Electrical Consumption				
Chiller	Calculated from Monitored Data	488	570	kW
Cooling Tower	Calculated from Monitored Data	39	25	kW
Primary Chilled Water Pump	Calculated from Spot Data	82	--	kW
Condenser Water Pump	Calculated from Spot Data	35	--	kW
CRAC Units (EACs) - Computer Room Cooling	Monitored	317	371	kW

The “HVAC Equipment Electrical” lists the measured electrical load for all cooling equipment that contributes to cooling the data center. This includes the chilled water plant components, the air handler units, and CRAC units. Chiller 3 consumes on average 796 kW. Chiller 4, which is the newest chiller, consumes 555 kW. The average efficiency of chiller 4 is superior to Chiller 3’s efficiency. This is discussed in more detail in the HVAC Efficiency Metrics and Energy Efficiency Recommendations Sections. The individual pump power is comparable to the total air handler power, and the cooling tower power. This presents a ripe opportunity for energy savings. The air handler power was monitored over a period of several days, and exhibited little variation during the monitored period. Finally, the CRAC unit power is substantial, and rivals the chiller power. Interestingly, the CRAC unit power did not vary much between weekend and weekday operation. Rumsey confirmed with the Building Operators that all but one CRAC unit were off in the areas serving the check processing. However, the CRAC unit

electrical loads show little decrease between weekday and weekend operation, averaging at 15%. This is consistent with the variation in the UPS input power, which was 10%.

The total chiller tonnage averaged at 860 tons for weekend periods. The CRAC units constituted a large part of this load, at 510 tons. The CRAC unit cooling load increased only slightly for the weekday measurement. On average, the CRAC cooling load showed a decrease of only 8% from weekend to weekday operation. This is consistent with the UPS input and CRAC unit power measurements.

Using the measured electrical consumption of the equipment is combined with the measured cooling loads to determine the portion of the HVAC equipment power accountable to the data center. This data is presented in more detail in the “Data Center Electricity End Use” Section.

LIGHTING

Lighting to the data centers is provided by T-8 tubular fluorescent lamps. Lighting to the data center floors alone could not be obtained. As a result, for purposes of the computing the data center end use, a lighting density of 1.5 W/sf is used. This assumption is used, since the facility's lighting resembles standard office lighting design.


DATA CENTER ELECTRICITY END USE

The measurements in the preceding sections are used to illustrate the Data Center Electricity End Use. The following table combines the HVAC, lighting, and computer power. The average energy use includes both weekend and weekday measurements. The same data is shown graphically for all measurement days.


TABLE 3. DATA CENTER AVERAGE ENERGY USE

	kW	Percent (%)
Computer Loads	1,420	47%
UPS Losses	355	12%
HVAC – Air Movement	353	12%
HVAC – Pumps & Chiller	747	25%
Lighting	119	4%
Total	2,993	100%

The data is also presented in graphical form below for each monitored day.



The above graph shows that the relative power consumption is fairly constant during the monitored period, with slighter lower computer loads on September 28 and 29. The average power during the monitored period, as a percentage, is shown in the graph below.



The largest data center energy consumer are the computer loads at 47% of the total. The UPS losses, though based on an efficiency estimate, are large even at an accuracy of $\pm 10\%$ (355 kW \pm 36 kW). The HVAC power consumes ~1100 kW, or 37% of the total data center power. This is a substantial amount of power, both in a relative, and absolute sense, and represents an opportunity for energy savings. The lighting loads are relatively small, compared to the other loads, at 119 kW, or 4% of total data center energy use; however, there is an opportunity to reduce lighting levels, and to implement lighting controls.

All the above areas present significant opportunities for energy savings. More details are included in the Energy Efficiency Recommendations Section. A more detailed discussion of efficiency metrics also follows.

WHOLE BUILDING ELECTRICITY END USE

Energy consumption of the whole building was obtained by taking spot readings of the instantaneous power (kW) from the utility meters. This data was consolidated with the monitored data to develop the end use for the whole building. The data is shown in two ways. The first method shows separates the data center electrical consumption from the office and miscellaneous spaces electrical consumption. The purpose is to illustrate the


relative contribution of the data center to the whole building power consumption. The table below lists the average whole building electricity consumption.

TABLE 4. WHOLE BUILDING AVERAGE ELECTRICITY END USE. METHOD 1.


Category	Average, Weekend		Weekday	
	kW	Percent (%)	kW	Percent (%)
Office & Facilities HVAC	423	9%	470	9%
Office Space (Lighting, Plug Loads) + Misc.	1422	30%	1547	31%
Data Center Computer Loads	1453	31%	1476	30%
UPS Losses	363	8%	369	7%
Data Center HVAC	970	20%	1022	20%
Data Center Lighting	111	2%	122	2%
Total Building Load	4742	100%	5007	100%

The whole building consumes an average of 5 MW of electricity. The weekday consumption is approximately 1 MW more than for the weekend, or 20% larger. This increase is seen mainly in the office space and miscellaneous loads, which include elevator loads, and smaller increases in other categories. In both modes of operation, the percentage breakdown of electricity use stays fairly constant. The major consumers are 1) the data center computer loads, 2) the office plug loads, lighting loads, and miscellaneous loads (which include elevator loads), and 3) the data center HVAC. They are approximately 30%, 30%, and 20%, respectively, or 70% together.

Measurements for the electrical consumption for each day are shown graphically below.



The above graph shows that the shift in electrical consumption between weekday and weekend operation is small. The graph below shows the overall average relative consumption in a different format.



The second method of showing the whole building electrical end use, is by the major categories of equipment, per their assignment to the utility meters. The distinction between data center and non-data center consumers is not made.


TABLE 5. WHOLE BUILDING AVERAGE ELECTRICITY END USE. METHOD 2.


Category	Average, Weekend		Weekday (Oct 1)	
	kW	Percent (%)	kW	Percent (%)
UPS Power	1816	38%	1932	33%
Chiller	555	12%	801	14%
Cooling Tower	69	1%	35	1%
Chiller Plant Pumps	251	5%	251	4%
EAC Power	338	7%	412	7%
All Other HVAC	180	4%	294	5%
Lighting, Office, Elevators, Misc.	1533	32%	2077	36%
Total Building Load	4742	100%	5802	100%

The largest consumer is the UPS power at approximately 1.9 MW, or 38% for weekend operation, and 33% for weekday operation. The difference between this value, and the one cited before is due to the UPS losses. The total HVAC power for the whole building

is approximately 30% for both scenarios, and the office, lighting, elevators and miscellaneous loads account for 34%. This representation further emphasizes the electrical consumption of HVAC equipment, and the relevance of energy efficiency measures. As shown earlier, approximately 2/3 of the HVAC power can be attributed to the data center cooling exclusively. The category of lighting, office, elevators, and miscellaneous sources emphasizes that lighting, as well as power management within offices is important, though this is not a focus of the study.

The relative electrical consumption is shown graphically below, for each day in the monitored period, and for the overall average.





The above graph shows three distinct areas: the UPS power, the lighting, office and miscellaneous loads, and the collection of consumers that make up the HVAC. Notice the chiller consumes approximately 12% of the total building power. Assuming a rough 5 MW of total consumption and \$0.12/kWh, this translates to a yearly cost of roughly \$0.6 million! The chiller plant pumps and CRAC /EAC power are also quite significant, at 5%, and 7%, respectively.

HVAC EFFICIENCY METRICS

The performance of the HVAC system can be evaluated based on energy efficiency metrics. Though the cooling power is represented in W/sf, a more useful metric for evaluating how efficiently the data center is cooled can be represented as a ratio of cooling power to computer power. This essentially removes the variable of how tightly packed the computers are.

The “theoretical cooling load” is the same as the sum of the computer loads, lighting loads, and fan energy. This is a good cross check. Differences can be attributed to the error in estimating UPS losses (which decrease the data center cooling load), duct losses, as well as unaccounted human load.¹³ The more traditional metrics of energy per ton of cooling (kW/Ton) are calculated for individual chillers, total chilled water plant (chillers, cooling towers, pumps) and the data center cooling efficiency. The data center cooling efficiency includes the chilled water plant power weighted by the data center load, data center air handler and CRAC unit power.

¹³ The fan energy included in the theoretical cooling load, removes an estimate of CRAC unit energy dedicated to the UPS that is cooled by the secondary loop.

TABLE 6. AVERAGE EFFICIENCY METRICS

Metric	Value	Units
Building Energy Density	3.6	W/sf
Data Center Computer Power Density	18	W/sf
Data Center Cooling Power Density	15	W/sf
Non Computer Load Density	21	W/sf
Cooling kW / Computer Load kW	0.7	--
Chiller 3 Efficiency	0.9	kW/ton
Chiller 4 Efficiency	0.6	kW/ton
Chilled Water Plant Efficiency	1.1	kW/ton
Total Data Center HVAC Efficiency	1.7	kW/ton
Theoretical Cooling Load *	572	Tons
Cooling Provided by CRAC Units	520	Tons
Cooling Provided by Air Handlers	121	Tons
Measured Cooling Load	641	Tons

* Based on computer loads, lighting loads, and fan energy.

The “cooling efficiency” is 0.7 Cooling kW/ Computer kW.¹⁴ This means that for 1 kW of energy input, only 1.4 kW of energy is removed. This value is slightly higher than the measured efficiencies of 0.5 kW/kW and 0.6 kW/kW at two other monitored data centers. The former utilizes air cooled chillers, and fan coil units, while the latter utilizes air cooled CRAC units. Both are considerably smaller in size, at 2500 sf, and 8600 sf. Another monitored site has an efficiency of 1.3 kW/kW. This data center utilizes a water cooled reciprocating chiller and computer room air handlers with humidification and reheat. This data center is also smaller at 8900 sf. Without going into details of these sites, it is interesting that a system that could operate efficiently (e.g., water cooled chilled water plant), isn’t necessarily more efficient than the standard air-cooled computer room air conditioners if the fundamental equipment, pumping, and air delivery systems are not efficient.

Chiller efficiencies were obtained for chillers 3 and 4, and on average are 0.9 and 0.6 kW/ton, respectively. This is expected, since chiller 3 was installed as part of the original plant, and chiller 4 was installed in 1994. However, though, Chiller 4 is a fairly new centrifugal chiller, its efficiency is not comparable to what it should be for the actual operating conditions, which are more favourable than the standard conditions at which the unit is rated at.¹⁵

¹⁴ The “Computer kW” includes the entire Exide UPS input, since the cooling kW is proportioned based on the secondary loop tonnage, which includes chilled water supplied to the Exide UPS CRAC units.

¹⁵ The rated conditions are: 80 °F entering condenser water temperature, chilled water setpoint of 42 °F. The operating conditions are: 68 °F entering condenser water temperature, chilled water setpoint of 42 °F.

The total chilled water plant efficiency, which includes the pumping and cooling tower power is 1.1 kW/ton. To put this in perspective, an efficient chilled water plant, such as all variable speed chilled water plant can operate at efficiencies between 0.5-0.7 kW/ton. The measured efficiency at this facility is comparable to standard chilled water plant efficiencies, with constant speed pumping, and constant speed chillers. These efficiencies typically vary between 0.8-1.2 kW/ton. The total HVAC efficiency, including air handling is 1.7 kW/ton. A standard HVAC design, utilizing similar equipment will typically operate at 1.5 kW/ton, while a range of 0.8-1.0 kW/ton is characteristics of efficient design.

Opportunities for energy savings exist, from simple measures to more complex and longer payback measures. These are described in the next section.

VI. Energy Efficiency Recommendations

IN HOUSE ENGINEERING RECOMMENDATIONS


Rumsey was provided with a report titled “Final Report, SFDC Energy Study”, dated October 5, 2001, by the facility’s in-house engineering group. We concur with many of the recommendations, which also addressed the whole building energy. In particular, the replacement, or addition of a dedicated VSD centrifugal chiller, and variable speed pumping on the secondary loop is suggested in this report. Note, however, we have not reviewed the calculation methodologies (as they were not apparent), or the assumptions utilized.

DEDICATED CHILLER FOR DATA CENTER, AND/OR OTHER CONTROL MEASURES

The facility currently utilizes the same chiller for the office cooling and data center cooling. The operating chiller cools water to a temperature of 42 °F, yet, the intercooler maintains a chilled water supply temperature of 48 °F for the CRAC units. The data taken in the monitored period, which includes 5 weekend days, and a 1 weekday, suggests that majority of the chiller load is absorbed by the CRAC unit load dedicated to the computer rooms (hereafter referred to as the CRAC unit load): During the weekend periods, the CRAC unit load averages at 59% of the chiller load, and was 58% of the chiller load on the weekday measurement. If the air handler load that is dedicated to the data center floors is included, then, the total measured data center load on weekend, and weekday operation is 73%, and 71%, respectively.

The chilled water setpoint has a direct correlation on the efficiency of the chiller. The measured efficiencies of chiller 3, and chiller 4 were 0.9 kW/ton, and 0.6 kW/ton, respectively. Greater efficiencies can be achieved if this supply temperature is raised. Though, each individual chiller is unique in terms of its operating characteristics, in general an “energy efficiency” rule of thumb is that the chiller’s efficiency increases by 2% for every 1 °F rise in chilled water setpoint. The graph below, based on measured data illustrates this point.

Comparison of Low Temperature and Medium Temperature Chillers




Suppose the Facility operated chiller 4 at a supply temperature of 48 °F, the chiller efficiency will be approximately 0.54 kW/Ton. If chiller 4 is operated instead of operating chiller 3 at a supply temperature of 42 °F, at the average data center load of 660 tons (average of weekend and weekday measurement), 2,081,376 kWh of energy will be saved per year, or approximately \$250,000 at \$0.15/kWh.

It is likely that the temperature is kept so low, because the chilled water flow to the farthest air handler is inadequate, and thus has to be compensated with cooler water. Investigating this problem could also be considered, and if it is possible for the main building air handlers to accept warmer chilled water, then this may be the simplest solution.

The most energy efficient solution is to install a data center dedicated chiller, preferably a centrifugal, variable speed drive (VSD) type. This allows the data center chiller to operate efficiently at part-load, and also allows for expansion in the data center, without concern of having cooling capacity. The existing chillers, including the newest chiller (4), are constant speed centrifugal type chillers. These chillers achieve lower loads, by adjusting inlet guide vanes of the compressor, which causes the refrigerant to swirl, and thus reduces the flow through the refrigerant cycle. A much more efficient mechanism of reducing refrigerant flow, and hence, chiller loading, is to reduce the speed of the motor, that is, install a VSD. The cost of VSDs in general has decreased dramatically in the past couple years, and VSD chillers are becoming more and more common. The graph below compares the typical efficiencies of different chiller types.

Recent conversations with the facility confirm that they are in the process of raising the chilled water setpoint to 48 °F for all chillers, and to 52 °F for the environmental chilled water. The facility is also in the process of installing VFDs on both Chillers 3 and 4. With these changes, the addition of a new, dedicated chiller is not required.

Comparison of Typical Chiller Efficiencies over Load Range



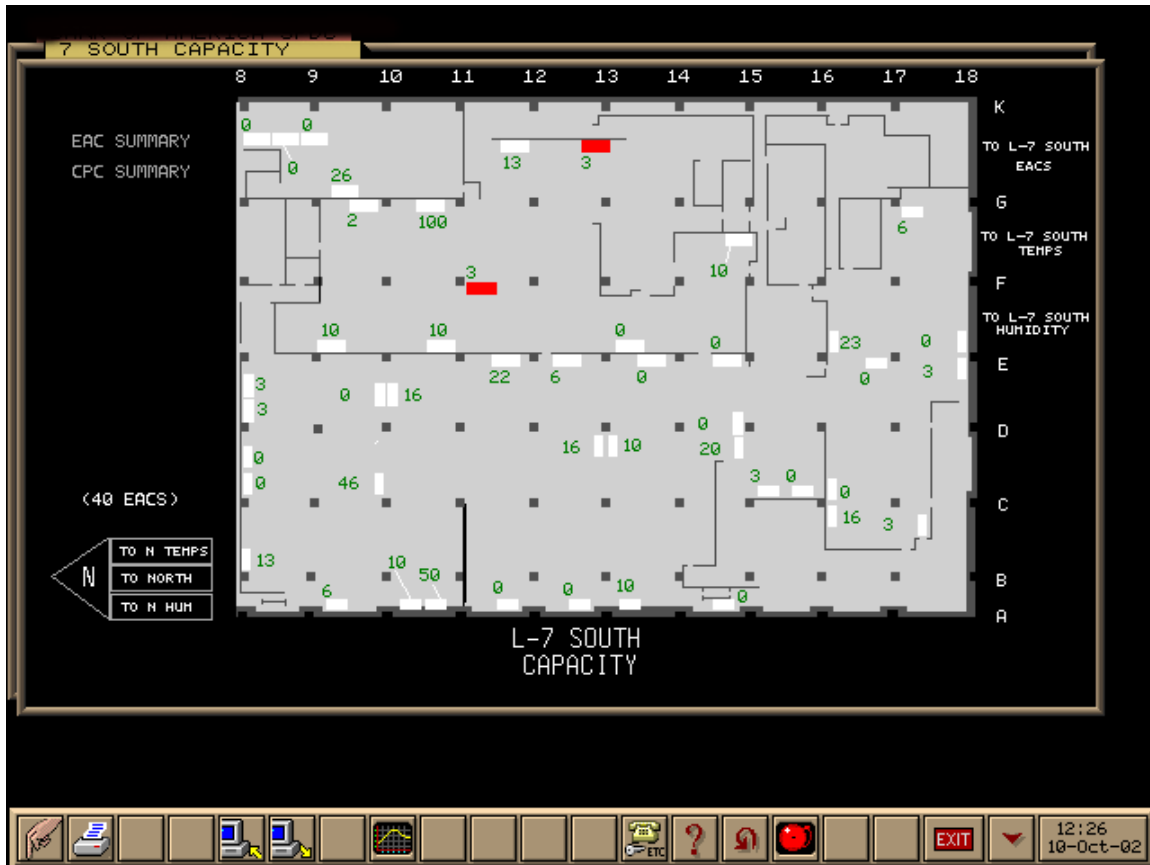
Chiller 1	250-Ton, Screw, Standard Efficiency, Air Cooled
Chiller 2	216 Ton, Screw, Water Cooled
Chiller 3	227-Ton, Centrifugal, Constant Speed, Water Cooled
Chiller 4	227-Ton, Centrifugal, Variable Speed, Water Cooled

The above graph clearly shows the advantage of a constant speed centrifugal VSD chiller over a constant speed centrifugal chiller at part-load conditions. At 35% loading, the VSD chiller is more than 40% more efficient than the constant speed centrifugal chiller. It is our understanding that the facility is contemplating a VSD retrofit of one of the older existing chillers. The facility should also consider a VSD retrofit on Chiller 4, since this chiller is the more efficient chiller.

CONVERSION TO VARIABLE SPEED PUMPING ON SECONDARY LOOPS AND OTHER PUMPING SAVINGS

A number of energy efficiency opportunities are available in the pumping of the chiller plant. The most obvious retrofit is to add a VFD to the secondary chilled water loop that serves the data center CRAC units. The system already consists of two-way valves on the

CRAC units, and maintains constant differential pressure in the chilled water supply line by controlling a bypass valve. Even at 48 °F supply temperature, the CRAC units' valves are barely open. The pictures below are snapshots of the CRAC unit monitoring system at the facility.



The picture shows that the CRAC units' valves are at varying positions, though, a majority are less than 20% open. Pumping savings are based on the cube law: pump power is reduced by the cube of the reduction in pump speed, which is directly proportional to the amount of fluid pumped. Assuming an average valve opening of 20%, for simplicity of calculations a linear acting valve, and the measured secondary pump power of 92 kW, a savings of 91 kW would result! This is a direct result of the cube law, and that the current pump is oversized for the existing CRAC units. These savings would amount to approximately 797,160 kWh, or roughly \$100,000/year. Currently, VFDs can be purchased at \$100/hp. The secondary pump's motor is 125 hp. A VFD would cost \$12,500, resulting in a simple payback of 1.5 months, or a return on investment (ROI) of 800%! The bypass valve should be permanently closed with this retrofit.

Premium efficiency motors and high efficiency pumps are recommended. During the retrofit, high efficiency motors were installed, which are more efficient than the motors installed originally. The existing motors should be retrofitted with premium efficiency

motors, and permitting, with VFDs. This would allow for a future retrofit to variable speed pumping on the primary chilled water side.

COOLING TOWER STAGING

These cooling towers have VFDs controlling the speed of the fans. VFDs are very energy efficient on cooling towers, because of the Cube Law savings discussed earlier. This allows the tower fans to modulate, based on varying outdoor air conditions, and to condenser water reset strategies. The facility's condenser water setpoint appears to be 68 °F. This is a fantastic way to achieve energy efficiency from the chillers. Just as chilled water setpoint affects chiller efficiency, lower condenser water temperatures reduce the chiller compressor work. However, the staging of the cooling towers is unclear. During some of the visits to the facility, it appeared that all eight towers were operating, as would be desired much of time. At other visits, it was observed that six towers were operating. The data also suggests that not all eight towers were operating during the monitored period, though, the operators had indicated they were. (The outdoor air conditions didn't change enough for the tower energy per fan to triple. This likely happened from the staging off of tower fans.) The sequencing information observed on the EMCS did not clearly indicate the staging sequence. It is recommended that the staging be based on the principal of operating all towers in parallel, with fewer towers as the fan speed reaches a minimum operating speed. This will ensure that cooling tower fan power is kept at a minimum.

REPLACEMENT / ELIMINATION OF INTERCOOLERS

Currently, the "intercooler" are shell and tube type heat exchangers. If the selection is made appropriately, a plate and frame heat exchangers will have a better approach for less pressure drop. The proper selection will save on pumping energy (if incorporated with the secondary pump VSD recommendation) and will allow the chilled water setpoint to be set to a higher value (currently at 42 °F), which will increase the efficiency of the chiller.

Based on the facility's plan of raising the chilled water setpoint for both building, and data center cooling, the intercooler can be completely eliminated. This unit adds additional pressure drop, and is not required if all systems can receive a common chilled water temperature.

ECONOMIZER BASED COOLING

A significant amount of cooling can be provided by outdoor air, particularly in this climate. Humidity control is often a concern in data center environments when outside air is introduced. This climate, is however, so moderate, that neither high humidity, nor low humidity is concern enough to not take advantage of outdoor air economizing. The air handlers that serve the data centers currently have fixed outside air dampers, and do not do economizing. It is encouraged that in this data center, and in future data centers in a similar climate, strongly consider using outdoor air economizing.

Recent conversations with the facility indicate that economizing should be possible with these air handlers, though they have not worked properly in the past. The facility has been working with its controls contractors to remedy this problem.

AIR MANAGEMENT

Humidity and Temperature Control

Over the past several years, the data center community has come to accept that tight humidity control is not an important factor for maintaining reliability of computers. Certainly, very low humidity can promote static electricity, however, tight humidity control, such as $50\% \pm 2\%$ is certainly unnecessary. It is encouraged, that in non-paper environments, the dehumidification be disabled (this promotes over-cooling, and re-heating) and the humidity control be broadened. (Observation of one of the CRAC units indicated a dead band of $50\% \pm 5\%$.) Currently, there are a collection of CRAC units that are permitted to do humidity control, and a number are not. This is in the right direction for saving energy.

Turn Off CRAC Units and Replace Perforated Tiles

The facility appears to be turning off selected CRAC units in areas that no longer have computers. This is certainly in the right direction, and can be pursued, so long as the underfloor is still adequately pressurized. As computers are moved, or removed, the placement of perforated tiles must also be managed. There are sand bags made for data centers that can be placed on perforated tiles.

Underfloor – Promote Thermal Stratification

The standard practice of cooling data centers employs an underfloor system fed by CRAC units. There are a number of potential problems with such systems: an underfloor system works on the basis of thermal stratification. This means that as the cool air is fed from the underfloor, it absorbs energy from the space, warming up as a result, and rises. In order to take advantage of thermal stratification, the return air must be collected at the ceiling level. CRAC units often have low return air grills, and are therefore, simply recirculating cool or moderately warmed air. Even if the grills are located on the top of the unit, the height of the CRAC units is unlikely to be high enough to capture warm air. Furthermore, they are often located along the perimeter of the building, and not dispersed throughout the floor area, where they can more effectively treat warm air. One alternative is to install transfer grills from the ceiling to the return grill.

Underfloor – Manage Cabling

Another common problem with underfloor supply is that the underfloor becomes congested with cabling, increasing the resistance to air flow. This results in an increase in fan energy use. A generous underfloor depth is essential for effective air distribution (We

have seen 3 feet in one facility. See www.nserc.gov). Also, it is essential that cabling be managed, and that when computers are moved, or removed, the associated cables are also removed.

Overhead System Alternative

An alternative to underfloor air distribution is high velocity overhead supply, combined with ceiling height return. This has been seen to work as efficiently as an underfloor system. A central air handling system can be a very efficient air distribution unit. Design considerations include using VFDs on the fans, low pressure drop filters, and coils.

Another common problem identified with CRAC units is that they are often fighting each other in order to maintain a constant humidity setpoint. Not only is a constant humidity setpoint unnecessary for preventing static electricity (the lower limit is more important), but it uses extra energy. A central air handling unit has a better ability to control overall humidity than distributed CRAC units.

Rack Configuration

Another factor that influences cooling in data centers is the server rack configuration. It is more logical for the aisles to be arranged such that servers' backs are facing each other, and servers' fronts are facing each other. This way, cool air is drawn in through the front, and hot air blown out the back. The Uptime Institute has published documents describing this method for air management.¹⁶ Our observations of the rack type areas of the data centers showed an inconsistent rack configuration. It is suggested that this arrangement be utilized in this data center, and for future data centers.

UPS REPLACEMENT

The UPS efficiency is likely to be poor, particularly, because the UPSs (Exide, and Teledyne) are loaded, on average at 54%, and 40%, respectively. The part load efficiency of the UPS drops dramatically, and observations at a 500 kVA UPS at a facility of comparable age to this facility, exhibited a UPS efficiency of 78%. This UPS was also partly-loaded. It is encouraged that when the new UPS is installed, efficiency is considered, and a gateway installed, so that the UPS can be monitored and trended at the EMCS.

Recent conversations with the facility indicate that both UPSs are operated such that each UPS should not be loaded by more than 50% for reliability purposes.

COMMISSIONING OF NEW SYSTEMS AND OPTIMIZED CONTROL STRATEGIES

Many times the predicted energy savings of new and retrofit projects are not fully realized. Often, this is due to poor and/or incomplete implementation of the energy

¹⁶ <http://www.upsite.com/TUIpages/whitepapers/tuiaisles.html>

efficiency recommendations. Commissioning is the process of ensuring that the building systems perform as they were intended to by the design. Effective commissioning actually begins at the design stage, such that the design strategy is critically reviewed. Either the design engineer can serve as the commissioning agent, or a third party commissioning agent can be hired. Commissioning differentiates from standard start-up testing in that it ensures systems function well relative to each other. In other words, it employs a systems approach.

Many of the problems identified in building systems are often associated with controls. A good controls scheme begins at the design well. In our experience, an effective controls design includes 1) a detailed points list, with accuracy levels, and sensor types, and 2) a detailed sequence of operations. Both of these components are essential for successfully implementing the recommended high efficiency chilled water system described above. Though commissioning is relatively new to the industry, various organizations have developed standards and guidelines. Such guidelines are available through organizations like the Portland Energy Conservation Inc., at www.peci.org, or ASHRAE, Guideline 1 - 1996.

OPTIMIZATION OF EMCS

The facility currently monitors chilled water flow at each chiller, and in the secondary loop, as well as temperatures. At a minimum, the chiller kW/ton, tonnage of the chiller, and in the secondary loop could be calculated and displayed on the EMCS. Though this is easy to calculate, having the value as a point and displaying it is valuable. If a correlation is made between the tower fan Hz, and power, (once the data for one point is made, then power is known at all speeds, since power is directly proportional to motor speed), then the variable portion of the electrical consumption by the entire chilled water plant is known. If the secondary pump is retrofitted with a VFD, then its consumption can also be added to this total kW/ton.

Another observation by the monitoring team is that the retrieval and transfer of Trend data to a remote computer can be performed easily using a copy and paste to clipboard function. Therefore, past data from the EMCS can be retrieved and analyzed on a continual basis quite easily, which will facilitate making controls and energy efficiency recommendations. Though the operations personnel view trend data and graphs on a continual basis *at* the EMCS, this functionality may help design engineers and non-operations personnel in making engineering decisions.

LIGHTING CONTROLS

The observed lighting appeared to be much larger than what is needed for data centers. In addition, all computer rooms that appeared to be unoccupied were fully illuminated. Lighting controls, such as occupancy sensors may be appropriate for areas that are infrequently, or irregularly occupied. If 24 hour lighting is desired for security reasons,

scarce lighting can be provided at all hours, with additional lighting being provided for occupied periods. The estimated lighting consumption for the data center is at 119 kW.¹⁷

¹⁷ Assuming 1.5 W/sf, and the data center gross areas, as obtained from facility's engineers.


APPENDICES – MONITORED DATA – FACILITY 7

CHILLED WATER PLANT GRAPHS


AIR HANDLER GRAPHS

CRAC UNIT / EAC GRAPHS

Facility B Data Center Chiller 4 Measured Power




Facility B Data Center Chiller 4 - Amps Sensor Error




Facility B Data Center


Chiller 4 Flow and Temperature - Oct 5 and 6




Facility B Data Center
Chiller 4 Power, Load and Efficiency - Oct 5 and 6




Facility B Data Center Total Chilled Water Plant Efficiency (Ch4) - Oct 5 and 6




Facility B Data Center Chiller 3 Measured Power




Facility B Data Center Chiller 3: Calculated Vs. Measured




Facility B Data Center Chiller 3 Flow and Temperature - Sep 28 and Sep 29




Facility B Data Center Chiller 3 Power, Load and Efficiency - Sep 28 and 29




Facility B Data Center Chiller 3 Power, Load and Efficiency - Sep 28 and 29




Facility B Data Center Total Chilled Water Plant Efficiency (Ch3) - Sep 28 and 29




Facility B Data Center Environmental Chilled Water Load - Sep 28, 29




Facility B Data Center
Environmental Chilled Water Load - Oct 5, 6, 7




Facility B Data Center Cooling Tower 1 Fan Power




Facility B Data Center Cooling Tower 5 Fan Power




Facility B Data Center Total Cooling Tower Fan Power




Facility B Data Center
AHU 4 CHW Supply & Return Temps




Facility B Data Center
AHU 4 Tonnage & CHW Flow




Facility B Data Center AHU 4 CHW Temps and Flow - Sep 28 Afternoon




Facility B Data Center
AHU 5 CHW Supply & Return Temps




Facility B Data Center AHU 5 Tonnage & CHW Flow



Facility B Data Center AHU 5 Fan Power Consumption




Facility B Data Center CRAC Power Consumption - Panel ATS 13



Facility B Data Center

CRAC Power Consumption - Panel ATS 13A



**Appendix IX. “Data Center Energy Benchmarking Case Study,
Facility 8”; PIER Project, Rumsey Engineers, Inc.; LBNL –
William Tschudi, 2003**

DATA CENTER ENERGY BENCHMARKING CASE STUDY

FEBRUARY 2003

FACILITY 8

SPONSORED BY:



LAWRENCE BERKELEY NATIONAL LABORATORY

PREPARED BY:



**99 LINDEN STREET
OAKLAND, CA 94607
(510) 663-2070**

Acknowledgements

Rumsey Engineers is grateful to the facility managers/directors, engineers and technicians for their generous assistance and cooperation. Special thanks to Christine Condon of PG&E for providing monitoring equipment on short notice. Thanks to the Lawrence Berkeley National Laboratory, and the California Energy Commission for funding this project.

Disclaimer

Neither Rumsey Engineers, LBNL nor any of its employees makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any data, information, method, product or process disclosed in this document, or represents that its use will not infringe any privately-owned rights, including but not, limited to, patents, trademarks, or copyrights.

Contents

I. EXECUTIVE SUMMARY	1
II. DEFINITIONS.....	9
III. INTRODUCTION.....	11
IV. SITE OVERVIEW.....	12
V. ENERGY USE – WHOLE BUILDING.....	12
VI. ENERGY USE – DATA CENTER 8.1.....	14
DATA CENTER 8.1: ELECTRICAL EQUIPMENT AND BACKUP POWER SYSTEM.....	14
DATA CENTER 8.1: COOLING SYSTEM.....	14
DATA CENTER 8.1: LIGHTING	16
DATA CENTER 8.1: SUMMARY OF MEASUREMENTS AND METRICS	16
VII. ENERGY USE – DATA CENTER 8.2.....	19
DATA CENTER 8.2: ELECTRICAL EQUIPMENT AND BACKUP POWER SYSTEM.....	19
DATA CENTER 8.2: COOLING SYSTEM.....	19
DATA CENTER 8.2: LIGHTING	21
DATA CENTER 8.2: SUMMARY OF MEASUREMENTS AND METRICS	21
VIII. ENERGY EFFICIENCY RECOMMENDATIONS.....	25
DESIGN COMMENTS.....	25
OPERATING COOLING TOWERS IN PARALLEL – INSTALL NEW NOZZLES.....	25
CONDENSER WATER TEMPERATURE RESET STRATEGY	25
CHILLED WATER RESET STRATEGY	26
CORRECT VFD GATEWAY PROGRAMMING AND CALIBRATE TURBINE FLOW METER.....	26
PERFORMANCE MONITORING	27
DATA CENTER 8.1 CRAC UNITS	27
UPS ROOM CRAC UNITS	28
DATA CENTER 8.1 UPS LOADING	28
TEMPERATURE CONTROLS	28

APPENDICES

A -- CHARTS OF MONITORED DATA

I. Executive Summary

Rumsey Engineers and the Lawrence Berkeley National Laboratory (LBNL) have teamed up to conduct an energy study as part of LBNL's Data Center Load Characterization and Roadmap Project, under sponsorship by the California Energy Commission (CEC). This study will aid designers to make better decisions about the design and construction of data centers in the near future. Data centers at four different organizations in Northern California were analyzed during the period of September 2002 to December 2002, with the particular aim of determining the end-use of electricity.

This report documents the findings for one of the case studies – termed Facility 8. Additional case study and benchmark results will be provided on LBNL's website (<http://datacenters.lbl.gov>), and are provided here for comparison purposes. The additional case studies are from this project, as well as a similar benchmarking study completed for the Pacific Gas and Electric Company (PG&E) in 2001.

Facility 8 contains two data centers, in one common building. Both data centers contain server type computers that resemble the server farms that became common as a result of the Internet Age.¹ The facility is an Internet Business Exchange center that serves as core hubs for critical IP networks and Internet operations. Data Center 8.1 was constructed prior to Data Center 8.2. Each has its own electrical and mechanical systems, and differs primarily in the type of cooling systems employed. Data Centers 8.1 and 8.2 are 26,200 square feet (sf), and 73,000 sf, respectively. The office spaces, and customer care areas are relatively small, and together, the data center areas total 75 % of the total building area.

Both data centers are supplied cool air through an overhead ducted system. Data Center 8.1 is cooled by air-cooled Computer Room Air Conditioning (CRAC) units. The CRAC units draw return air in the front, and supply air to the ducted system through the top of the units. The condenser units are located on the roof. Data Center 8.2 is cooled by a central chilled water plant, consisting of water cooled centrifugal, variable speed (VSD) chillers, cooling towers, and a primary / secondary pumping system. The chilled water is supplied to several large, centralized air handlers, which supply cool air to the overhead ducted system. Note, Facility 8 does not utilize a raised floor cooling system.


The current computer energy loads are listed in the table below. A qualitative estimate of the loading of the racks was made, and the future computer energy loads were estimated based on this loading. (The occupancy estimate is based on the facility's self assessment of occupancy.) For comparison purposes the computer loads of other data centers studied in this project, and a previous PG&E project are also included. The computer loads are also shown graphically.

¹ Based on the rack configuration, high density of computers, and absence of the large mainframe servers that were common in older data centers.

CURRENT AND FUTURE COMPUTER LOADS

Data Center	Data Center Area (sf)	Computer Load (kW)	Computer Load Energy Density (W/sf)	Occupancy (%)	Projected Computer Load Energy Density (W/sf)
Data Center 8.1	26,200	222	8	30%	27
Data Center 8.2	73,000	1,059	15	30%	50
Data Center 6.1	2,400	155	65	80%	81
Data Center 6.2	2,501	119	48	50%	95
Data Center 7	116,700	1,395	12	82%	14
Data Center 1	62,870	1,500	24	75%	32
Data Center 2	60,400	2,040	34	65%	52
Data Center 3	25,000	1,110	44	85%	52

COMPUTER LOADS




The measured computer load densities at Data Center 8.1 and 8.2 are relatively small, compared to some of the data centers studied. They are 8.5 W/sf and 15 W/sf, respectively. At an estimated occupancy level of 30%, the full occupancy computer load densities are projected at 28 and 50 W/sf, respectively. Facility 8 resembles the environments of the data centers from the previous utility study, yet the load densities are

smaller. This could be due to the fact that the previous utility study occurred during the peak of the dot com frenzy. (late 2000).

WHOLE BUILDING ENERGY USE

The whole building electricity end use is shown in the figure below. The whole building consumes an average of 2.6 MW of electricity. The major consumers are, in ranking order 1) the data center computer loads, 2) the data center HVAC loads, 3) the data center UPS losses, 4) other/miscellaneous loads which include HVAC, lighting and plug loads for office spaces, front lobby and equipment rooms. The transformer losses, at 3% only account for the step-down in electricity from 12 kV to 480 V. Data center lighting power consumption is relatively small at 1% of the total building's load.²




² Note, it is not possible to make efficiency observations for either of the data centers, based on the whole building data alone.

DATA CENTER 8.1 ENERGY USE

The remaining energy loads of Data Centers 8.1 and 8.2 include HVAC equipment power, lighting, and uninterruptible power supply (UPS) efficiencies. The measurements are shown in graphical format below and tabular format in the report.

DATA CENTER 8.1 ENERGY BALANCE



A large percentage, approximately 54%, of the total electrical load is from the HVAC equipment, which is exclusively the CRAC unit power. This value is very large, and in previous studies has been less than the server or computer loads. Lighting is a very small percentage of the total energy, at 2%. The UPS losses are 6%. This value is likely to decrease as the data center becomes more fully loaded, since the UPS efficiency improves at higher loads. Data Center 8.1 is wired such that each PDU is served by either Side A or Side B, where each side contains three UPS's. Therefore, at the PDU level, the redundancy is n+1, but at the UPS level, n+2.

The performance of the HVAC system can be evaluated based on energy efficiency metrics. Though the cooling power is represented in W/sf, another interesting metric for evaluating how efficiently the data center is cooled can be represented as a ratio of cooling power to computer power.

DATA CENTER 8.1 EFFICIENCY METRICS

Metric	Value	Units
--------	-------	-------

Metric	Value	Units
Data Center Computer Power Density	8.5	W/sf
Data Center Cooling Power Density	12	W/sf
Cooling kW : Computer Load kW	1.4	--

The data center computer load density is small, relative to what is observed at other data centers. Hence, the cooling energy density is also small, at 12 W/sf, relative to other facilities. The “cooling efficiency”, which is the efficiency normalized to the computer power is 1.4 Cooling kW/Computer kW. This means that for every 1 kW of heat generated by the computers, 1.4 kW of power is required to remove it. This “cooling efficiency” is at the high end of the range, but is not surprising since air cooled CRAC units are utilized. For comparison purposes, the cooling efficiencies measured at other facilities, and their HVAC system types are listed below.


COOLING EFFICIENCY AT OTHER DATA CENTERS

System Type	Efficiency (kW/kW)
Air cooled chillers and fan coil units	0.5
Air cooled CRAC units	0.8
Water cooled reciprocating chiller and CRAC units	1.5
Water cooled centrifugal chiller plant with CRAC units	0.7

DATA CENTER 8.2 ENERGY USE


The electrical end use for Data Center 8.2 is shown in two graphs below. The first gives an overall breakdown of energy consumption based on major categories, such as computer/server loads, HVAC energy, UPS losses, and lighting.

DATA CENTER 8.2 ENERGY BALANCE



The energy use breakdown for Data Center 8.2 differs from what is observed for Data Center 8.1 significantly. The server or computer loads constitute a much larger percentage, at 63%. UPS losses, at 13% are larger, since these UPS's are lightly loaded. HVAC energy is relatively much less at 23%. This highlights that the cooling systems are operating more efficiently in Data Center 8.2, than in Data Center 8.1. The air movement energy, namely from the air handlers, constitutes 41% of the total HVAC energy, and the chilled water plant 59%. The HVAC energy is broken down further in the next graph.

DATA CENTER 8.2 ENERGY BALANCE WITH HVAC BREAKDOWN



Majority of the chilled water plant energy is from the chiller at 64%, pumping at 29%, and cooling tower fan energy at 7%.

The efficiency metrics for Data Center 8.2 are listed in the table below. Note, they include the common metrics kW/Ton for chillers, total chilled water plant (chillers, cooling towers, pumps), and the data center cooling efficiency.

DATA CENTER 8.2 EFFICIENCY METRICS

Metric	Value	Units
Data Center Computer Power Load Density	14.5	W/sf
Data Center Cooling Power Load Density	5.3	W/sf
Cooling kW : Computer Load kW	0.4	--
Chiller Efficiency	0.4	kW/ton
Chilled Water Plant Efficiency	0.6	kW/ton

The cooling power load density is 5.3 W/sf. The “cooling efficiency” is 0.4 Cooling kW/Computer kW. This means that for 1 kW of energy added by the computer loads, 0.4 kW of energy is required to remove it. This value is the most efficient among all the data centers studied. This is not surprising since the HVAC equipment is high efficiency. The

chiller efficiency is 0.4 kW/ton, chilled water plant efficiency is 0.6 kW/ton, which is in the range of what is seen in efficient chilled water plants.

In general, the cooling system for Data Center 8.2 is significantly more efficient than the cooling system for Data Center 8.1, by 71%. However, there are additional energy efficiency and optimization opportunities for Data Center 8.2 , as well as Data Center 8.1, which are discussed in the “Energy Efficiency Recommendations” section of the report.

II. Definitions

Data Center Facility	A facility that contains both central communications equipment, and data storage and processing equipment (servers) associated with a concentration of data cables. Can be used interchangeably with Server Farm Facility.
Server Farm Facility	A facility that contains both central communications equipment, and data storage and processing equipment associated with a concentration of data cables. Can be used interchangeably with Data Center Facility. Also defined as a common physical space on the Data Center Floor where server equipment is located (i.e. server farm).
Data Center Floor / Space	Total footprint area of controlled access space devoted to company/customer equipment. Includes aisle ways, caged space, cooling units, electrical panels, fire suppression equipment, and other support equipment. Per the Uptime Institute Definitions, this gross floor space is what is typically used by facility engineers in calculating a computer load density (W/sf). ³
Data Center Occupancy	This is based on a qualitative estimate on how physically loaded the data centers are.
Data Center Cooling	Electrical power devoted to cooling equipment for the Data Center Floor space
Data Center Server/Computer Load	Electrical power devoted to equipment on the Data Center Floor. Typically the power measured upstream of power distribution units or panels. Includes servers, switches, routers, storage equipment, monitors, and other equipment.
Computer/Server Load Measured Energy Density	Ratio of actual measured Data Center Server Load in Watts (W) to the square foot area (ft ² or sf) of Data Center Floor. Includes vacant space in floor area.

³ Users look at watts per square foot in a different way. With an entire room full of communication and computer equipment, they are not so much concerned with the power density associated with a specific footprint or floor tile, but with larger areas and perhaps even the entire room. Facilities engineers typically take the actual UPS power output consumed by computer hardware and communication equipment in the room being studied (but not including air handlers, lights, etc.) and divide it by the gross floor space in the room. The gross space of a room will typically include a lot of areas not consuming UPS power such as access aisles, white areas where no computer equipment is installed yet, and space for site infrastructure equipment like Power Distribution Units (PDU) and air handlers. The resulting gross watts per square foot (watt/ft2-gross) or gross watts per square meter (watt/m2-gross) will be significantly lower than the watts per footprint measured by a hardware manufacturer in a laboratory setting.

Computer /Server Load Projected Energy Density	Ratio of forecasted Data Center Server Load in Watts (W) to square foot area (ft ² or sf) of the Data Center Floor if the Data Center Floor were fully occupied. The Data Center Server Load is inflated by the percentage of currently occupied space.
Cooling Load Tons	A unit used to measure the amount of cooling being done. Equivalent to 12,000 British Thermal Units (BTU) per hour.
Chiller Efficiency	The power used (kW), per ton of cooling produced by the chiller.
Cooling Load Density	The amount of cooling (tons) in a given area (ft ² or sf).

III. Introduction

This report describes the measurement methodology and results. The facility contained two data centers, which were measured independently. In each data center, electricity end use was determined. This means that the energy consumed by all equipment related to the data center was measured. Such equipment includes the actual computer power consumption, the data center air conditioning equipment, the lighting, and the inefficiencies associated with the uninterruptable power supply (UPS). The computer load density is also determined based on the gross area of the data center. This number, in watts per square foot (W/sf) is the metric typically used by facility engineers to represent the power density. Based on a qualitative observation of the data center occupancy, the computer load density at full occupancy is extrapolated.

Additional information was collected so that the efficiencies of the cooling equipment could be calculated. These efficiencies are compared to the design efficiencies. Opportunities for energy efficiency improvements are described, which are based on observation of the mechanical system design, and measured performance.

IV. Site Overview

Facility 8 is located in the Silicon Valley of California. Facility 8 contains two data centers located in one common 131,700 square feet (sf) building. The data centers, hereafter referred to as Data Center 8.1 and Data Center 8.2, are 26,200 sf and 73,000 sf, respectively. Data Center 8.1 was constructed prior to Data Center 8.2. The data centers are adjacent to one another and do not have a physical wall of separation. Each has its own electrical and mechanical systems, and differs only in the type of cooling systems employed. The office spaces, and customer care areas are relatively small, and together, the data center areas total 75 % of the total building area. Both data centers house servers and storage drives, and operate 24 hours a day.

V. Energy Use – Whole Building


The 131,700 sf building is made up of 99,200 sf of data center space, 23,300 sf of equipment room space and 9200 sf of office space.

TABLE 1. WHOLE BUILDING SUMMARY AND METRICS

	Measurement	Units	
Data Center Computer Loads	1281	kW	50%
Data Center UPS Losses	246	kW	10%
Data Center HVAC Loads	694	kW	27%
Data Center Lighting	37	kW	1%
Transformer Losses	87	kW	3%
Other / Miscellaneous	225	kW	9%
Whole Building Power Consumption	2571	kW	100%
Whole Building Load Density	20	W/sf	

The data centers' computer loads are the highest consumer at 50%, followed by the total data centers' HVAC loads at 27%. The data center power consumption altogether accounts for 88% of the whole building load. Transformer losses are only accounted for the transformers that step down the power from 12 kV to 480 V. The "Other / Miscellaneous" category include: equipment rooms and office space plug, HVAC and lighting loads, and other miscellaneous loads. The data center load density shown in the table is based on the square footage of the whole building. The table above is also represented by the pie chart below.

FIGURE 1. WHOLE BUILDING ENERGY BALANCE



VI. Energy Use – Data Center 8.1

DATA CENTER 8.1: ELECTRICAL EQUIPMENT AND BACKUP POWER SYSTEM

The facility utilizes six 375 kVA Liebert uninterruptible power supplies (UPS) to provide a constant supply of power to the data center at constant delivery voltage (480/277 V). The UPS's convert AC current and stores it as DC current in multiple battery packs. When constant voltage is needed, it is converted back to AC current. Each PDU can be served by two UPS's, therefore, at the PDU level, the redundancy is n+1.

In the event of a power loss, three 750 kW diesel generators will provide backup power.

Trended power measurements were taken at the UPS in order to determine computer plug loads.

TABLE 2. UPS ELECTRICAL MEASUREMENTS

	Electrical Use ¹	Units
UPS A-1 Output	30.4	kW
UPS A-2 Output	30.0	kW
UPS A-3 Output	30.7	kW
UPS B-1 Output	42.2	kW
UPS B-2 Output	44.6	kW
UPS B-3 Output	44.6	kW

¹ Trended Data from 11/25/02 – 11/26/02.

DATA CENTER 8.1: COOLING SYSTEM

The data center is cooled with eighteen (18) ducted, air-cooled, constant volume CRAC units. Air is delivered through the top of the unit to overhead ducting and returned through grills on the front of the unit. Note that this data center does not have a raised floor. The nominal capacity of the CRAC units is 18 tons. The computer room air conditioners have constant-speed fans designed to deliver air at a maximum of 10,200 cubic feet per minute (cfm). Electric humidification and reheat are also utilized by these units, though, the site has disabled humidity control in some of the CRAC units. A physical separation wall does not exist, which allows for ventilation air exchange between Data Centers 8.1 and 8.2.



FIGURE 2. DATA CENTER CRAC UNIT

During the monitored period, temperature was maintained at an average of 64 °F and relative humidity was maintained between 45% and 55% for those with enabled humidity control.⁴

Spot measurements of power consumption were taken for the CRAC units, which includes the condenser fan power (denoted as CU-XX). Spot measurements of some individual units are listed in the table below.

TABLE 3. COOLING EQUIPMENT ELECTRICAL AND LOAD MEASUREMENTS

Equipment	Spot / Monitored	Date	Measurement	Units
AC-12	Spot	11/26/02	32.2	kW
CU-12	Spot	11/26/02	2.8	kW
AC-14	Spot	11/26/02	6.4	kW
CU-14	Spot	11/26/02	0 (Off)	kW
AC-20	Spot	11/26/02	17.2	kW
CU-20	Spot	11/26/02	2.6	kW

1 Condenser unit, CU-14 was off during measurement due to cool outside air conditions.

⁴ During the walk thru on November 26, 2002, active humidity control was observed on CRAC Units 6 and 21.

DATA CENTER 8.1: LIGHTING

Lighting in the data center was unusual when compared with the other data centers in this study; specifically this data center was significantly darker than other data centers where lighting was more similar to regular office lighting. Only lights illuminating the walkways were permanently on. The equipment cages had overhead fluorescent lights that turned on only when occupied. An estimate of the lighting density is made based on qualitative observation of the lighting levels.

Lighting Power Density: 0.5 W/sf

Lighting Power: 13 kW

DATA CENTER 8.1: SUMMARY OF MEASUREMENTS AND METRICS

The table below summarizes the electrical measurements for the data center equipment.


TABLE 4. SUMMARY OF ELECTRICAL MEASUREMENTS

Computer Loads	222	kW	38%
UPS Losses ¹	33	kW	6%
HVAC	309	kW	54%
Lighting	13	kW	2%
Total Energy Use	578	kW	100%

¹ Estimated at 15% of the computer load, based on recorded efficiencies for UPS's in Data Center 8.2.

These results are shown graphically in the pie chart below.

FIGURE 3. DATA CENTER 8.1 ENERGY BALANCE



The computer loads, based on the UPS power supply is 38% of the data center energy usage. Air conditioning energy is the largest consumer at 54%. This is due to the fact that most, or all CRAC units are operating, though the data center is not fully loaded. Losses in the UPS system account for 6% of the data center energy consumption. These losses are more than the lighting, which amount to only 2% of total energy use.

The electrical and cooling loads can be represented by different metrics. The most commonly used metric among mission critical facilities is the computer load density in watts consumed per square foot. However, the square footage is not always consistent between designers. Some data centers use kVA/rack or kW/rack as a design parameter. Our definition of “Data Center Floor Area” includes the gross area of the data center, which includes rack spaces, aisle spaces, and areas that may eventually contain computer equipment. Per the Uptime Institute, the resulting computer load density in watts per square foot is consistent with what facility engineers use, though this is different from the “footprint” energy density that manufacturers use. We have also calculated the W/sf based on the rack area alone. In addition to the previous metrics, the “non-computer” energy densities are calculated, based on the “data center area”. Using the data center occupancy⁵ the computer load density at 100% occupancy is projected.

⁵ A qualitative assessment of how physically full the data center is. In this facility, occupancy was determined by a visual inspection of how full the racks in place were.

TABLE 5. ELECTRICAL CONSUMPTION METRICS

Data Center Gross Area ¹	26,200	sf
"Occupied" %	30%	Estimated from visual inspection
<i>Based on Gross Area:</i>		
Computer Load Density	8	W/sf
Non-Computer Load Energy Density	14	W/sf
Projected Computer Load Density	27	W/sf

¹ Gross area includes spaces between racks; does not include entire building area.

The computer load density based on the data center area (gross area) is 8 W/sf. At full occupancy, the computer load density is projected to be 27 W/sf. The non-computer energy density, which includes HVAC, lighting, and UPS losses, is measured at 14 W/sf. Both the computer and non-computer load densities are relatively small compared to other facilities studied.

A useful metric for evaluating how efficiently the data center is cooled can be represented as a ratio of cooling power to computer power. Although, there is a small amount of human activity, the computer energy load dominates.

TABLE 6. HVAC EFFICIENCY METRICS

Metric	Value	Units
Cooling kW: Computer Load kW	1.4	--

From the above table it is shown that the “cooling efficiency” is 1.4 kW/kW. This suggests that for every 1 kW of heat generated, 1.4 kW is consumed to remove the heat. This “cooling efficiency” is at the high end of the range, based on measurements from the other facilities in this study. Other facilities measured efficiencies of 0.5 kW/kW and 0.6 kW/kW, which utilize air cooled chillers and fan coil units, and air cooled CRAC units, respectively. Another monitored site has an efficiency of 1.3 kW/kW, which utilizes a water cooled reciprocating chiller and CRAC units with humidification and reheat. Though a water cooled chiller plant could operate extremely efficiently, it will not if the fundamental equipment, air delivery and pumping systems are inefficient.

VII. Energy Use – Data Center 8.2

DATA CENTER 8.2: ELECTRICAL EQUIPMENT AND BACKUP POWER SYSTEM

The facility utilizes five Liebert 1100kVA uninterruptible power supplies to provide a constant supply of power of constant delivery voltage to the data center. The UPS converts AC current and stores it as DC current in multiple battery packs. When the voltage is needed, it is converted back to AC current. In the event of a power loss, two 2000 kW diesel generators will provide power.

Here as well, trended measurements were taken over a period of two days⁶, as well as spot power measurements at the UPS display panel. The spot measurements were necessary in order to obtain both input, and output power, and hence the UPS efficiencies. The spot measurements deviated from the averaged trend data by less than 1 kW, and are reported in the table below.

TABLE 7. UPS ELECTRICAL MEASUREMENTS

	Electrical Use ¹	Units
UPS-3 Input	117	kW
UPS-3 Output	79	KW
UPS-3 Efficiency	68	%
UPS-4 Input	362	kW
UPS-4 Output	314	KW
UPS-4 Efficiency	87	%
UPS-5 Input	205	kW
UPS-5 Output	166	KW
UPS-5 Efficiency	81	%
UPS-6 Input	290	kW
UPS-6 Output	248	KW
UPS-6 Efficiency	86	%
UPS-7 Input	298	kW
UPS-7 Output	252	KW
UPS-7 Efficiency	85	%

¹ Spot data taken at 3:47 pm, 11/26/02.

The UPS efficiencies are poor, which is likely due to the low loads. UPS's become more efficient as their load increases.

DATA CENTER 8.2: COOLING SYSTEM

The data center is cooled by a chilled water system. The chilled water system consists of four identical 500 ton, Carrier VFD centrifugal chillers. The nominal efficiency of the

⁶ 11/25/02 – 11/26/02.

chillers is 0.56 kW/Ton.⁷ The chillers are piped in a primary-secondary loop configuration, and two are typically operating at a time. The primary chilled water pumps are constant-speed, with 15 hp motors; secondary chilled water pumps are variable-speed, with 40 hp motors; and condenser water pumps are constant-speed with 20 hp motors. The cooling towers have variable-speed fans with 25 hp motors. There are four cooling towers. Variable-speed chillers are more efficient at part-load conditions, and the staging strategy reflects this fact: a chiller is brought on-line when the operating chiller load exceeds 80%, and a chiller is taken off-line when the operating load goes below 40%. During the monitored period two chillers, primary chilled water pumps, two secondary chilled water pumps, and two cooling towers were operating.

The chilled water feeds thirteen (13) large air handling units (AHUs), nine (9) water-cooled CRAC units, and one (1) small AHU. The thirteen AHUs serve the data center primarily. The CRAC units and the small AHU provide cooling to the electrical and mechanical equipment rooms, including the UPS rooms.

The thirteen air handling units that serve the data center are identical and have variable-speed fans and two-way cooling coil valves. The cooling capacity is 105 tons each at a design airflow of 56,000 cfm. All air handlers are equipped with outdoor air economizers, and take full advantage of this feature. Humidity is monitored by five humidity sensors. When the humidity in the space goes below this limit, the air handlers recirculate return air. Temperature is monitored in the zones. The air delivery volume is varied by opening and closing a branch duct damper, such that the worst case “zone” is satisfied.

Power consumption of various HVAC equipment, chilled water flow, and chilled water temperatures were obtained using the building’s control system over a period of two days. The chiller control panel, and variable speed drive (VSD) control panels had gateways to the EMCS, which facilitated retrieving power data. Power consumption of the chiller and variable speed driven equipment were obtained from the EMCS in this manner. This did not include the constant speed pump data, such as the primary chilled water pump data, and condenser pump data.

Verification of the certain data points was performed where the readings seemed erroneous. This included verification of chiller, secondary pump, and cooling tower fan data. The VFD programming was found to be off by a factor of 10, which indicated an unreasonable amount of cooling tower fan and secondary pump power! The originally obtained data was corrected by using the verification data obtained on the follow up site visit.⁸ In addition, the chilled water flow rate was verified, as this measurement is critical to calculate an accurate chiller efficiency.⁹

⁷ Based on 1000 gpm, entering and leaving chilled water temperatures of 56 °F, and 44 °F, respectively, and entering condenser water temperature of 80 °F.

⁸ The VFD gateways for the secondary pumps, and cooling tower fans were reporting power consumption that was 10 x the actual power consumption.

⁹ It was verified that the secondary pump frequency (Hz) did not vary between the original monitored date, and the follow-up site visit. Therefore, the flow rate verified using the Controlotron Ultrasonic flow meter is used to calculate the chiller load.

The spot measurements and average of trended measurements are listed in the table below. Please refer to the Appendix for graphs of the measurements over the entire monitored period. The chiller pump and chiller power are proportioned to the data center cooling load in order to properly determine the electrical end use in the data center.

TABLE 8. COOLING EQUIPMENT ELECTRICAL LOAD MEASUREMENTS

Equipment	Spot / Monitored	Date	Measurement	Units
Chiller – Total	Monitored	11/26/02	164	kW
Chiller - Proportioned by Data Center Load	Spot	11/26/02	154	kW
Primary Chilled Water Pumps - Total	Spot	11/26/02	23.0	kW
Primary Chilled Water Pumps - Proportioned by Data Center Load	Spot	11/26/02	21.7	kW
Secondary Chilled Water Pumps – Total	Spot	11/26/02	9.6	kW
Secondary Chilled Water Pumps - Proportioned by Data Center Load	Spot	11/26/02	9.0	kW
Cooling Tower – Total	Spot	11/26/02	13	kW
Cooling Tower - Proportioned by Data Center Load	Spot	11/26/02	12.3	kW
Condenser Water Pumps – Total	Spot	11/26/02	32.3	kW
Condenser Water Pumps - Proportioned by Data Center Load	Spot	11/26/02	30.4	kW
AHU-10	Monitored	11/26/02	12.3	kW
AHU-11	Monitored	11/26/02	12.2	kW
AHU-12	Monitored	11/26/02	11.8	kW
AHU-13	Monitored	11/26/02	12.0	kW

DATA CENTER 8.2: LIGHTING

The lighting strategy in the data center was even more efficient than the lighting in Data Center 8.1, and thus, lighting power density is estimated at a lower level.

Lighting Power Density: 0.3 W/sf

Lighting Power: 24 kW

DATA CENTER 8.2: SUMMARY OF MEASUREMENTS AND METRICS

The table below summarizes all of the equipment electrical measurements for the data center.


TABLE 9. SUMMARY OF ELECTRICAL MEASUREMENTS

Computer Loads	1059	kW	62%
UPS Losses	213	kW	13%

HVAC - Air Movement	157	kW	9%
HVAC – Chilled Water Plant	228	kW	14%
Lighting	24	kW	1%
Total Energy Use	1681	kW	100%

These results are shown graphically in the pie chart below.


FIGURE 4. DATA CENTER 8.2 ENERGY BALANCE



The energy use breakdown differs from what was observed in Data Center 8.1. The computer loads, based on the measured UPS power supply amounts to 63% of the data center energy usage. Cooling and pumping energy is the second largest consumer at 13 %. Power consumption for air movement from the AHUs is 9%. Together, the HVAC component accounts for 23% of data center energy use, which is significantly much lower than what was characterized in Data Center 8.1. This indicates that Data Center 8.2's HVAC system is operating more efficiently. Losses at the UPS total 13% of the data center energy use. Similar to Data Center 8.1, the UPS's were operating inefficiently due to the low loads. The percentage of lighting power consumption was less for this data center, estimated at 1%.

The HVAC consumption by components is shown below. The data center was maintained at an average temperature of 64 °F through the monitored period.

FIGURE 5. HVAC ENERGY BREAKDOWN



This graph shows that within the chilled water plant, the pumping accounts for 4% of the total data center power, cooling towers 1%, and chillers 9%. Majority of the chilled water plant energy is from the chiller at 64%, pumping at 29%, and cooling tower fan energy at 7%. The pumping energy is significant, specifically because the primary chilled water and condenser water pumps are constant speed. Together they consume 55 kW, which is 85% of the total pumping energy of 65 kW. The cooling tower energy is low, though, further savings may be possible by running more fans in parallel. Conversations with the facility suggest that this was attempted, but adequate distribution of water on the evaporative media was not achieved, and therefore, the energy increased. This is discussed in the Energy Efficiency Recommendations Section.

Commensurate with the discussion under Data Center 8.1, metrics are calculated for the data center energy use, and energy efficiency. To briefly reiterate, the computer load density is based on both gross area, which we equate to “data center floor area”, and on rack floor area. The density is also extrapolated to 100% occupancy to predict future loads.

TABLE 10. ELECTRICAL CONSUMPTION METRICS

Data Center Gross Area	73,000	sf
"Occupied" %	30%	Estimated from visual inspection
<i>Based on Gross Area:</i>		
Computer Load Density	15	W/sf
Non-Computer Load Energy Density	9	W/sf
Projected Computer Load Density	50	W/sf

The computer load density based on the data center area (gross area) is 14.5 W/sf. At full occupancy, the computer load density is projected to be 50 W/sf. The non-computer energy density, which includes HVAC, lighting, and UPS losses, is measured at 9 W/sf.

Commensurate with Data Center 8.1, the energy efficiency metrics are shown in the table below.

TABLE 11. HVAC EFFICIENCY METRICS

Metric	Value	Units
Cooling kW: Computer Load kW	0.4	--
Theoretical Cooling Load	345	Tons
Cooling Provided by Chilled Water Plant to Data Center	391	Tons
Chiller 1 & 2 Design Efficiency ¹⁰	0.56	kW/ton
Chiller 1 & 2 Combined Efficiency	0.4	kW/ton
Chilled Water Plant Efficiency	0.6	kW/ton

From the above table it is shown that the “cooling efficiency” of approximately 0.4 kW/kW is significantly smaller than the “cooling efficiency” for Data Center 8.1, by 71%. This means that for the same amount of heat generated, Data Center 8.2 uses 71% less energy to remove that load. This can be explained by the differences in equipment; a few water-cooled variable-speed chillers with large variable-speed air handling units is superior to several small, self contained, constant-speed air cooled CRAC units.

The performance of the chillers is efficient at 0.4 kW/ton. This is better than the ARI rated efficiency, which is expected since the operating condenser water temperature is lower than the rating condenser water temperature. The total chilled water plant efficiency is 0.6 kW/ton, which includes the chillers, pumps and cooling towers. This is also in the efficient range of chilled water plant efficiencies.

¹⁰ The nominal efficiencies cannot be directly compared to the average operating efficiencies, since the nominal efficiencies are based on full load capacities, and the specific conditions cited previously.

VIII. Energy Efficiency Recommendations

DESIGN COMMENTS

The designs employed in Data Center 8.1 and Data Center 8.2 are vastly different. Overall, the design approach in Data Center 8.2 incorporates energy efficiency design concepts, and the energy efficiency metrics obtained from the data monitoring confirms that the design and operation is indeed more efficient.

Such design principles include a central chilled water plant that utilizes centrifugal VFD chillers, variable speed pumping, VFD cooling tower fans, condenser water temperature reset strategies, central air handlers, no or little humidity control, outside air economizing, gravity damper exhaust (without return fans), high ceiling and return. In addition to the presence of efficient equipment, extensive monitoring is done by the EMCS, and features such as “gateway” from the chiller control panel, and VSD control panels are present. Furthermore, the EMCS software has block programming that is easily accessible from the interface, allowing for real time adjustment of control variables, such as valve tuning parameters.

The next sections describe our energy efficiency recommendations, which are directed at optimizing the existing equipment, and control strategies.

OPERATING COOLING TOWERS IN PARALLEL – INSTALL NEW NOZZLES

Currently, the facility is unable to operate all cooling towers in parallel. In theory, it uses less energy to operate all towers in parallel because of the cube law. The cube law states that fan power is reduced by the cube of the reduction in fan speed, which is directly proportional to the amount of air moved. Therefore, if the speed of the fan is reduced by half, then the power will be reduced by a factor of eight. The facility has apparently experimented with operating cooling towers in parallel, but saw a decrease in efficiency because the water was not evenly distributed on the fill material. We recommend installing smaller nozzles so that more even flow is achieved.

CONDENSER WATER TEMPERATURE RESET STRATEGY

Currently, the condenser water temperature is reset based on the compressor mapping, per the manufacturer’s recommendation. However, a more aggressive reset strategy may be considered. One approach may be to reset the condenser water temperature based on an accurate reading of the outdoor wet-bulb temperature, with an appropriate differential for the approach. This control method has been successfully implemented at other locations, though a good quality sensor is necessary for its success.¹¹

¹¹ The Oakland Museum is such an example.

Another would be to reset the condenser water temperature based on maintaining a fixed, or minimum differential refrigerant pressure. This is possible because the chillers have gateways, which means that the refrigerant pressures are available at the EMCS. Such strategies are supported by chiller manufacturers such as Trane.¹² These options should be discussed with the manufacturer, and coordinated with the controls and chiller vendors before implementation.¹³

CHILLED WATER RESET STRATEGY

The monitored data shows a chilled water temperature of 46 °F. However, majority of the air handler valves are only 20% open. This means that the chilled water temperature could be raised to improve the efficiency of the chiller (at a small sacrifice of secondary pumping energy). In addition to increasing the efficiency of the chiller, the higher chilled water temperature will result in less dehumidification at the coil. This will reduce the load on the chiller, and result in additional energy savings. Per conversation with the facility, this strategy will be tested and implemented with an incremental approach.

Currently, the condenser water temperature mapping is based on a standard fixed chilled water temperature (probably 45 °F). The chiller compressor will have to be “reset”, such that it is “trained” to operate efficiently at the new operating conditions. Previous experience has shown chillers with sophisticated computer algorithms do not operate efficiently at new operating conditions, if the machine is not trained properly. We recommend working with the chiller manufacturer to reset the chiller controls for new operating conditions, if necessary.

CORRECT VFD GATEWAY PROGRAMMING AND CALIBRATE TURBINE FLOW METER

This site is unique in that the variable speed drives have gateways installed, which allows direct communication of all data with the EMCS. Thus, information such as power consumption of the secondary chilled water pumps, cooling tower fans, air handler supply fans are easily available from the EMCS! Power data was obtained by this method, however, as noted in the report, the programming of the VFD gateways appears to be incorrect, and should be corrected. This involves performing power measurements at the equipment, and simultaneously viewing the networked power data at the EMCS. Our validation indicates that the gateway is reading a power value that exceeds the actual power consumption by 10 times.

As described in the report, the chilled water flow rate was verified, and found to be less than the reading of the turbine flow meter. These meters are prone to error, and it is

¹² See Trane Engineering Bulletin CTV-PRB006-EN

¹³ More sophisticated reset strategies are achievable by adjusting all chilled water plant setpoints in real time via monitoring of power measurements. Though this is possible with the monitoring capabilities at this site, the complexity is not justified for a semi-variable speed chilled water plant.

suggested that they be calibrated on a yearly, or biyearly basis. Though this is not relevant to efficiency, it ensures that the monitored chiller efficiency data is accurate.

PERFORMANCE MONITORING

For the purpose of viewing the efficiency of the system in real time, we recommend that the following points be created, and displayed on the building's energy management control system: chiller kW/Ton, and chilled water plant kW/Ton (to include VFD kW, in addition to chiller kW). This will enable the facility to easily track the effect that controls changes have on the chiller plant efficiency.

DATA CENTER 8.1 CRAC UNITS

The "cooling efficiency" in Data Center 8.1 was 1.5 kW/kW. Again, this data center is one of the worst performers in the entire study. The CRAC units are equipped with humidity control, and in general, maintain relative humidity between a range of 45-55%. Humidity control utilizes extra energy, and is also an additional maintenance headache. Often, the CRAC units will be fighting each other, that is, some will be humidifying and some will be dehumidifying. In fact, this scenario was observed at this facility: one CRAC unit had a relative humidity of 39%, and was humidifying, while an adjacent CRAC unit had a relative humidity of 53%, and was dehumidifying. And, yet, another adjacent CRAC unit had a relative humidity of 43% and was erroneously in dehumidification mode! The net effect is poor humidity control, and wasted energy. More and more, it is becoming accepted in the data center community that a broad range of humidity is acceptable, and the fluctuation of humidity may be more important than maintaining humidity within a tight dead-band. Therefore, it is recommended that the facility investigate this range, and subsequently disable humidity control on all or some of the CRAC units.¹⁴

The CRAC units were designed for a fully loaded data center and do not have the ability to reduce its air flow delivery to better meet the current load conditions. Turning off some of the CRAC units can be an effective way to reduce fan energy. In turn, this will create a higher temperature difference between the air supply and return paths resulting in a more efficient refrigeration cycle. Conversations with the facility indicate that this was attempted, but that the reduced pressure drop resulted in the fans running off the fan curve. One method of mitigating this problem is to replace the existing fan sheave with a larger fan sheave, such that a slower fan speed is achieved. This effectively downsizes the fan, such that it will operate in the desired range of the fan curve. It is also recommended that the CRAC units be equipped with back-draft dampers, so that backwards air flow is prevented when a CRAC unit is turned off.

¹⁴ Conversations with the facility suggest that prior maintenance personnel may have attempted to eliminate humidity control in some of the CRAC units. This is strongly encouraged.

UPS ROOM CRAC UNITS

CRAC units with overhead ducting are utilized to cool the UPS rooms. In one of the UPS rooms, several supply registers are within 10 feet to the CRAC units. Capping off the registers closest to the CRAC units will minimize short circuiting of the cold supply air. Short circuiting occurs when the cold supply air is drawn back into the return path of the air handler before exchanging any heat energy. Although, the cooling coil will not see a load from the short circuited air (besides fan motor heat), fan energy that has been used to move the air is wasted.

DATA CENTER 8.1 UPS LOADING


Data Center 8.1 is structured, such that two groups of UPS's serve the computer loads. Each PDU can be served by either group of UPS's. Within the group of UPS's, there are three UPS's, therefore, an n+2 redundancy. At low-load conditions, we recommend that the facility consider operating the UPS's, such that one UPS is off-line on each side. This still preserves the n+1 redundancy, and if UPS's are rotated on a regular basis, the charge should be maintained.

TEMPERATURE CONTROLS


The temperatures in the data center are observed to be 64-66 °F. Such cold temperatures are not necessary. It is recommended that the temperature setpoints be increased to a more moderate temperature, such as 70 - 72 °F. It is beginning to be recognized that the *fluctuation* in environmental conditions adversely affects computer performance more than the actual humidity and temperature setpoints.

APPENDICES – MONITORED DATA – FACILITY 8, DATA CENTER 8.1

Facility 8 Data Center 8.1 Total Power




Facility 8 Data Center 8.1 UPS Power




APPENDICES – MONITORED DATA – FACILITY 8, DATA CENTER 8.2


Facility 8 Data Center 8.2 Total Power




Facility 8 Data Center 8.2 UPS Power




Facility 8 Data Center 8.2 Chiller 1 Temperatures



Facility 8 Data Center 8.2 Chiller 1 Power




Facility 8 Data Center 8.2 Chiller 2 Temperatures




Facility 8 Data Center 8.2 Chiller 2 Power




Facility 8 Data Center 8.2 Secondary Chilled Water Loop Temperatures




Facility 8 Data Center 8.2 Secondary Chilled Water Loop Flow




Facility 8 Data Center 8.2 Total Chiller Power




Facility 8 Data Center 8.2 Chilled Water Plant Tonnage & Efficiency



Facility 8 Data Center 8.2 Cooling Tower Power




Facility 8 Data Center 8.2 Secondary Chilled Water Pump Power



Facility 8 Data Center 8.2

Air Handler Fan Power



**Appendix X. “Data Center Energy Benchmarking Case Study,
Facility 9”; PIER Project, Rumsey Engineers, Inc.; LBNL –
William Tschudi, 2003**

DATA CENTER ENERGY BENCHMARKING CASE STUDY

FEBRUARY 2003



FACILITY 9

Sponsored By:



**Lawrence Berkeley National Laboratory
Environmental Energy Technologies Division**

Prepared By:



**99 Linden Street
Oakland, CA 94607
(510) 663-2070**

Acknowledgements

Rumsey Engineers is grateful to the facility managers/directors, engineers and technicians for their generous assistance and cooperation. Special thanks to Christine Condon of PG&E for providing monitoring equipment on short notice. Thanks also to the Lawrence Berkeley National Laboratory and the U.S. Department of Energy.

Disclaimer

Neither LBNL nor any of its employees makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any data, information, method, product or process disclosed in this document, or represents that its use will not infringe any privately-owned rights, including but not, limited to, patents, trademarks, or copyrights.

Contents

I. EXECUTIVE SUMMARY	1
II. DEFINITIONS	3
III. SITE OVERVIEW	5
IV. ENERGY USE.....	9
UNINTERRUPTIBLE POWER SUPPLY	10
DATA EQUIPMENT.....	11
COOLING SYSTEM	16
LIGHTING	27
SUMMARY MEASUREMENTS AND METRICS	27
V. ENERGY EFFICIENCY RECOMMENDATIONS.....	34
GENERAL GUIDELINES FOR FUTURE CONSTRUCTION.....	34
<i>Efficient Chilled Water System</i>	<i>34</i>
<i>Air Management.....</i>	<i>36</i>
<i>Air Management – Rack Configuration.....</i>	<i>36</i>
<i>Commissioning of New Systems and Optimized Control Strategies.....</i>	<i>36</i>
<i>Lighting Controls.....</i>	<i>37</i>
SPECIFIC RECOMMENDATIONS	37
<i>Revisit Humidity Control Strategy.....</i>	<i>37</i>
<i>Run Fewer CRAC Units.....</i>	<i>37</i>
<i>Rearrange Perforated Floor Tiles for More Effective Cooling.....</i>	<i>38</i>
<i>Reduce Lighting</i>	<i>38</i>

APPENDICES

- A – CHARTS OF MEASURED DATA
- B – MEASUREMENT OF UPS EFFICIENCY
- C – COOLING PROVIDED BY AIR HANDLER AH17


I. Executive Summary

Rumsey Engineers and the Lawrence Berkeley National Laboratory (LBNL) have teamed up to conduct an energy study as part of LBNL's Data Center Load Characterization Project. This study is intended to help designers make better decisions about the design and construction of data centers in the near future. This report describes the analysis of a data center in Sacramento, California. Measurements were conducted on-site from January 6 to 8, 2003, with the particular aim of determining the end-use of electricity. The identity of the organization that owns this data center is kept anonymous. The facility that houses the data center is referred to throughout this report as Facility 9.

The report begins with definitions of some basic terms. This is followed by an overview of the data center and the cooling system that serves it.

Section IV, Energy Use, is the main focus of the report. It is sub-divided by equipment type – UPS, data equipment, cooling system, and lighting. The electric power consumed by the equipment, and notable equipment behavior, is examined in detail. The section concludes with summaries and metrics. The following chart is an example from the summary.

Figure 13. Data Center Energy Balance



The report concludes with recommendations for saving energy at Facility 9, some general and some specific. The two chilled water plants examined in this study are due for removal in early 2003; Facility 9 will obtain chilled water from a new central plant that will serve several buildings. Given these circumstances, the specific recommendations in this report focus on the data center itself, and not the existing chilled water plants. The specific recommendations address humidity control, effective computer cooling, and lighting.

II. Definitions

Air Flow Density	The air flow (CFM) in a given area (ft ² or sf)
Chiller Efficiency	The power used (kW), per ton of cooling produced by the chiller.
Computer /Server Load Projected Energy Density	Ratio of forecasted Data Center Server Load in Watts (W) to square foot area (ft ² or sf) of the Data Center Floor if the Data Center Floor were fully occupied. The Data Center Server Load is inflated by the percentage of currently occupied space.
Computer Load Density – Rack Footprint	Measured Data Center Server Load in Watts (W) divided by the total area that the racks occupy, or the rack “footprint”.
Computer Load Density per Rack	Ratio of actual measured Data Center Server Load in Watts (W) per rack. This is the average density per rack.
Computer/Server Load Measured Energy Density	Ratio of actual measured Data Center Server Load in Watts (W) to the square foot area (ft ² or sf) of Data Center Floor. Includes vacant space in floor area
Cooling Load Density	The amount of cooling (tons) in a given area (ft ² or sf)
Cooling Load Tons	A unit used to measure the amount of cooling being done. Equivalent to 12,000 British Thermal Units (BTU) per hour.
Critical Load	Electrical load of equipment that must keep running in the event of a power failure. Such loads are typically served by an Uninterruptible Power Supply (UPS), that uses a bank of batteries to support the load when the normal source of power fails. The batteries can support the load for only a short period. In some facilities the equipment is shut down gracefully and turned off until normal power returns. In other facilities a backup generator, typically diesel-powered, comes on-line and provides power for a longer period of time.
Data Center Cooling	Electrical power devoted to cooling equipment for the Data Center Floor space
Data Center Facility	A facility that contains data storage and processing equipment (servers) associated with a concentration of data cables.

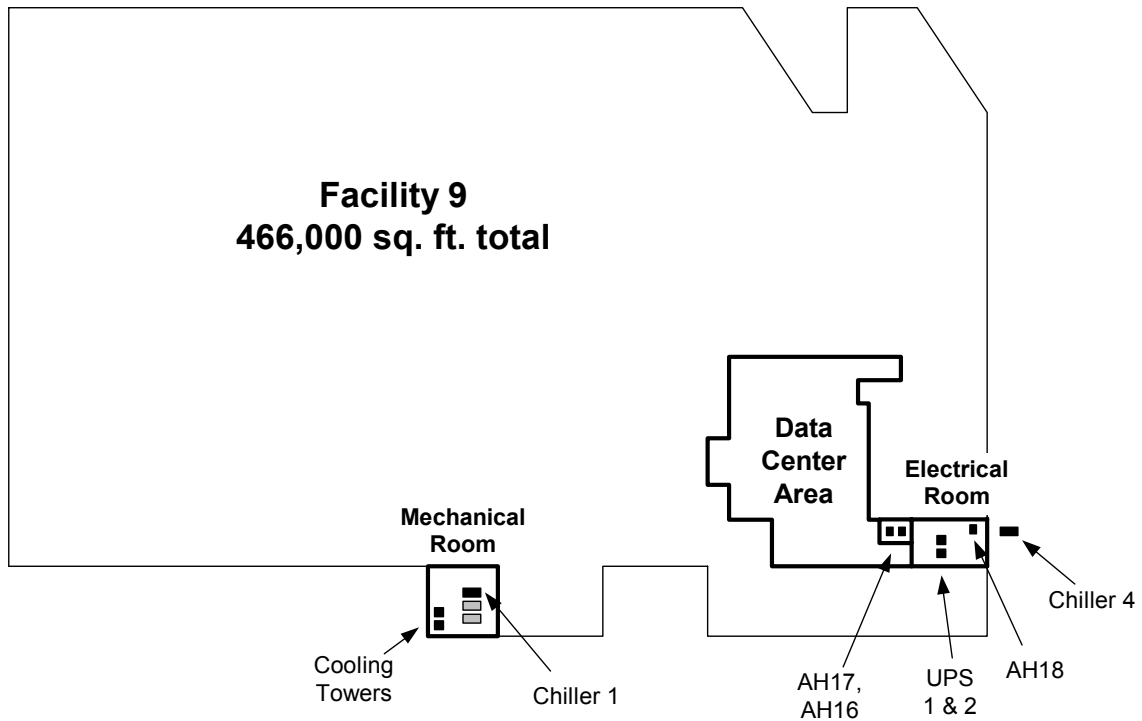
Data Center Floor / Space	Total footprint area of controlled access space devoted to company/customer equipment. Includes aisles, caged space, cooling units, electrical panels, fire suppression equipment, and other support equipment. Per the Uptime Institute Definitions, this gross floor space is what is typically used by facility engineers in calculating a computer load density (W/sf). ¹
Data Center Occupancy	This is based on a qualitative estimate on how physically loaded the data centers are.
Data Center Server/Computer Load	Electrical power devoted to equipment on the Data Center Floor. Typically the power measured upstream of power distribution units or panels. Includes servers, switches, routers, storage equipment, monitors, and other equipment.

¹ Users look at watts per square foot in a different way. With an entire room full of communication and computer equipment, they are not so much concerned with the power density associated with a specific footprint or floor tile, but with larger areas and perhaps even the entire room. Facilities engineers typically take the actual UPS power output consumed by computer hardware and communication equipment in the room being studied (but not including air handlers, lights, etc.) and divide it by the gross floor space in the room. The gross space of a room will typically include a lot of areas not consuming UPS power such as access aisles, white areas where no computer equipment is installed yet, and space for site infrastructure equipment like Power Distribution Units (PDU) and air handlers. The resulting gross watts per square foot (watt/ft2-gross) or gross watts per square meter (watt/m2-gross) will be significantly lower than the watts per footprint measured by a hardware manufacturer in a laboratory setting.

III. Site Overview

All of Facility 9 is housed in a single, 466,000 square-foot building; see Figure 1.

Figure 1. Facility 9 Site Plan



On their web site (<http://datacenters.lbl.gov/What.html>), LBNL defines “data center” as follows:

“We define a data center as a special facility that performs one or more of the following functions:

- Store, manage, process, and exchange digital data and information;
- Provide application services or management for various data processing, such as web hosting internet, intranet, telecommunication and information technology.

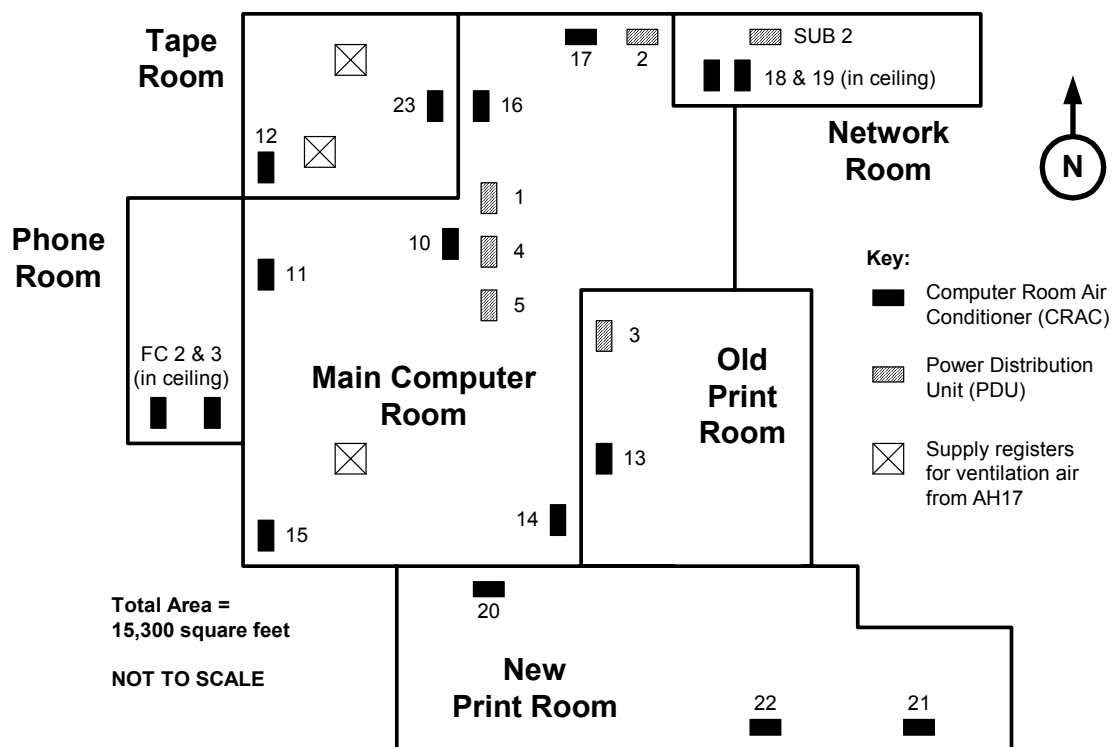
We do not consider spaces that primarily house office computers including individual servers associated with work stations as data centers.”

For this report, we treat the entire area shown in Figure 2 as a single data center. It comprises about 3% of the total building area. This data center is the focus of this report.

The data center contains six rooms. The Main Computer Room houses the facility's primary data processing servers. The Network Room is dedicated to the building's computer network servers. The Old and New Print Rooms contain large, continuous-feed printers. Six robotic memory-tape silos are installed in the Tape Room, and the telephone switching equipment is found in the Phone Room.

All of the computer and printer equipment loads in the data center are served by six Power Distribution Units (PDUs) located in the data center. All the PDUs are served by a pair of uninterruptible power supply (UPS) systems located in the electrical room.

Figure 2. Data Center Floor Plan



The data center is cooled primarily by twelve floor-mounted Computer Room Air Conditioning (CRAC) units.² In addition, there are two fan coil units and two more CRAC units mounted in the interstitial space above the ceiling tiles. The CRAC units are capable of humidity control; the fan coils are not. Makeup air for the data center is provided by air handler AH17. All fourteen CRAC units, the two fan coils, and air handler AH17 use chilled water for cooling. All seventeen units serve only the data center.

² Facility 9 uses the term CAC instead of CRAC. This report uses the term CRAC, to be consistent with reports on other data center facilities.

The Main Computer Room, Tape Room, and the Old Print Room have raised floors, consisting of 2-foot by 2-foot tiles approximately 1 foot above the underlying slab. Most of the tiles are solid, but approximately 10% of them are perforated and distributed across the floor area.

In these raised floor areas, the CRAC units supply air directly to the underfloor space. Air rises through perforated tiles, and returns through grills located in the top of the CRAC units.

In the areas without a raised floor, the CRAC units deliver air upwards. The air is ducted up to the space above the ceiling tiles, and then the duct branches to several ceiling supply registers. Likewise, the fan coil units installed in the ceiling deliver their air via ducts and ceiling supply registers.

Air handler AH17 is housed in a small room between the data center and the electrical room. The makeup air from this unit is delivered to three ceiling supply registers – one in the Main Computer Room, and two in the Tape Room.

All of the floor-mounted CRAC units are capable of humidity control. The ceiling-mounted units do not have humidity control. Humidification is provided by electric humidifiers, and reheat for the dehumidification process is also electrical.

Table 1 lists the basic specifications of all the CRAC and fan coil units.

Table 1. CRAC and Fan Coil Unit Descriptions

Label	Room	Make	Nominal Cooling Capacity (tons)	Has Digital Display	Floor/Ceiling Model	Air Flow
CRAC-10	Main Computer	Liebert	20	yes	Floor	Down
CRAC-11	Main Computer	Liebert	20	yes	Floor	Down
CRAC-14	Main Computer	Liebert	20	yes	Floor	Down
CRAC-15	Main Computer	Liebert	25	yes	Floor	Down
CRAC-16	Main Computer	Liebert	25	yes	Floor	Down
CRAC-17	Main Computer	Liebert	25	yes	Floor	Down
CRAC-18	Network	Liebert	5	no	Ceiling	Side
CRAC-19	Network	Liebert	5	no	Ceiling	Side
CRAC-13	Old Print	Liebert	20	yes	Floor	Down
CRAC-20	New Print	Liebert	5	yes	Floor	Up
CRAC-21	New Print	Liebert	5	yes	Floor	Up
CRAC-22	New Print	Liebert	20	yes	Floor	Up
CRAC-12	Tape	Liebert	10	yes	Floor	Down
CRAC-23	Tape	Liebert	10	yes	Floor	Down
FC-2	Phone	?	?	no	Ceiling	Side
FC-3	Phone	?	?	no	Ceiling	Side

Chilled water is distributed to all the CRAC and fan coil units, and to the make-up air handlers, by a chilled water pipe loop. Chilled water is introduced to the loop from two different chillers. See “Cooling System” under Section IV of this report for a description of the chilled water plants.

The entire HVAC system is controlled and monitored by a Barber Coleman Network 8000 system.

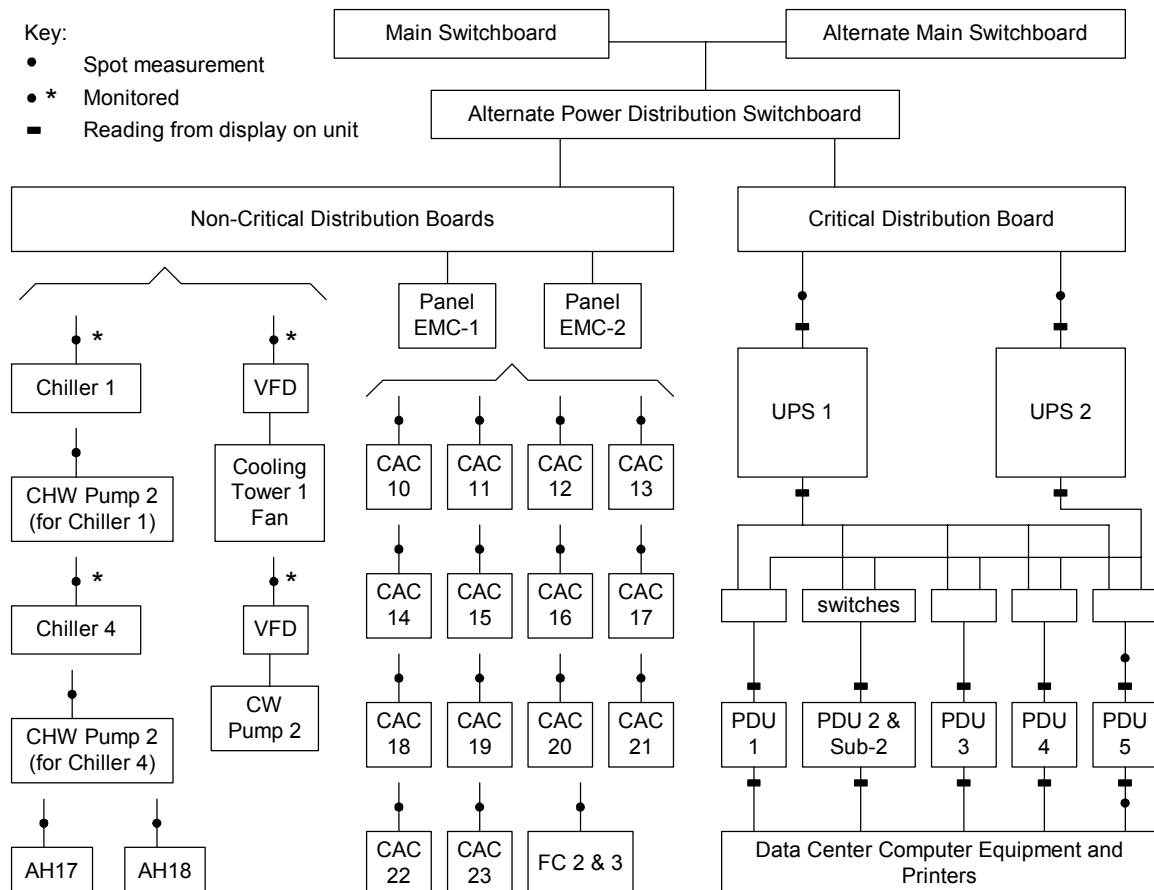
Both the main chilled water plant and the air-cooled chiller were due for removal shortly after the measurements for this study were completed. A new building is being added to the site. All the buildings on the site will be served by a new, central chilled water plant.

IV. Energy Use

Data Center Facility 9 is served by two utility substations. If one substation fails, the other substation picks up the load. If both substations fail, the UPS system supports the critical load for about 100 minutes, permitting a graceful shutdown of the computer equipment. There is no on-site power generator.

All of the computer and printer equipment in the data center is considered critical load, and is served by the UPS system via the PDUs. The HVAC equipment is considered non-critical load. Figure 3 is a simplified electrical diagram that shows the points of measurement. A Powersight PS3000 clamp-on power meter was used for all spot measurements and for longer periods of monitoring. The UPS units and the PDUs are equipped with digital display panels that indicate input and output load.

Figure 3. Electrical Measurement Points



UNINTERRUPTIBLE POWER SUPPLY

Two 500 kVA Liebert model UDA63500A36A uninterruptible power supply (UPS) units provide power to the critical loads in the data center. The UPS converts alternating current to direct current and charges a battery bank. Direct current from the batteries is converted back to alternating current and is fed to the data centers.

Based on the current total data center load, either UPS has sufficient capacity to serve the load by itself. Both UPS units operate continuously, however; this is required to keep both battery banks charged. The power supplied to and from the UPS was measured to determine how much of their capacity is being used, and how efficiently they are operating.

Table 2. UPS Electrical Measurements³

	Units	UPS 1	UPS 2	Combined
Input	kW	227.1	35.8	262.8
Output	kW	203.5	17.0	220.5
Loss	kW	23.6	18.8	42.3
Efficiency	%	89.6	47.3	83.9
Load Factor	%	40.7	3.4	16.8

Figure 4. UPS 2



³ Input and output values are spot readings of the UPS on 1/7/03. See Appendix B for details of the measurement process.

The measurements show that UPS 1 has a higher load factor than UPS 2 (40.7% vs. 3.4%), and is operating much more efficiently (89.6% vs 47.3%). The data center manager stated that this is a temporary condition; normally the UPS systems are loaded equally.

DATA EQUIPMENT

The data center consists of six rooms. The Main Computer Room contains server racks, tape drives, tape shelves, and several workstations. The computer equipment is arranged close together in the north end of the room, and less so in the remainder of the room. See Figure 5.

Figure 5. Main Computer Room



The Network Room contains two rows of server racks, that nearly fill the entire room. The Old and New Print Rooms contain large, automated printers. The Old Print Room is approximately 80% full of equipment, and the New Print Room is about 30% full. See Figures 6 & 7.

Figure 6. Old Print Room



Figure 7. New Print Room



The Tape Room contains six large, vertical cylinders that house robotic tape drives. There is no space to add additional units of this type. A photo of this room appears on the cover of this report. Finally, the Phone Room is full of racks of phone switch gear.

All computer equipment in the data center receives power from one of the five PDUs. The PDUs receive power from the UPS and remove spikes and transients.

Figure 8. PDU-4



All the PDUs have a 125 kVA nominal capacity. All of the PDUs are equipped with a digital display. To ascertain the accuracy of these displays, the input of PDU-5 was measured directly. PDU-5 was chosen over the others as it was the most accessible of the units. The results are shown in Table 3.

Table 3. PDU-5 Measurements

	Time of Day, 1/7/03	Source of Data	Power	Power Factor	Effic.	Load Factor	Loss
			kW	--	%	%	kW
Input	17:21	Powersight PS3000	19.49	0.96	94.8	13.0	1.0
Output	17:37		18.48	0.97			
Input	17:40	PDU Display Panel	n/a	n/a	n/a	n/a	n/a
Output	17:40		18	1.00			

Even though PDU-5 was only 13% loaded at the time of measurement, it was running at 94.8% efficiency. High efficiencies are typical of this type of power conditioning equipment. The display panel readings show close agreement with the directly measured values.

To obtain the total load served by the PDUs, all the PDU display panels were read within a 14-minute period on 1/8/03 as shown in Table 4.

Table 4. PDU Readings⁴

Label	Output Power (kW)	Power Factor	Load Factor (%)	Time
PDU-1	30	1.0	24	11:06
PDU-2	65	1.0	52	11:09
PDU-3	47	1.0	38	11:16
PDU-4	25	1.0	20	11:04
PDU-5	19	1.0	15	11:02
Combined	186	--	30	--

The total electrical load served by the PDUs was 186 kW at the time of measurement. This represents 30% of the total nominal capacity of the PDUs.

As a check on the total PDU load, the output of the UPS units was read from the UPS display panels less than 2 hours later. The readings are shown in Table 5.

⁴ All readings taken on 1/8/03 from PDU displays. PDU-2 is actually two units: PDU-2, and PDU-Sub-2. The latter serves the Network Room, and does not have a display panel. The display panel on PDU-2 reports the total power for both units.

Table 5. UPS Readings⁵

Label	Output Power (kW)	Load Factor (%)	Time
UPS 1	173.5	35	12:58
UPS 2	16.5	3	12:58
Combined	190.0	19	--

The total output of the UPS units is in good agreement with the PDU readings. Note that the total UPS output is less than was observed on the previous day, as shown in Table 2. The measurement team observed that the UPS output fluctuated somewhat during the period of measurement; this is likely due to the large printers cycling on and off.

⁵ Readings taken on 1/8/03.

COOLING SYSTEM

CHILLED WATER PLANTS

The CRAC units, fan coils, and three air handlers receive chilled water from a distribution loop; see Figure 11.

Chilled water is introduced to the distribution loop from two sources. The first is the main chilled water plant consisting of three, water-cooled centrifugal chillers. Only Chiller 1 ran during the period of measurement; it is a Carrier 150-ton unit. According to the maintenance crew, Chillers 2 and 3 do not run during the winter months. The maintenance crew says that all of the chilled water produced by Chiller 1 during the monitored period served only the data center. These statements are in keeping with observations made by the measurement team. As discussed below, the office space served by air handler AH16 received all of its cooling via air-side economizing, and did not require any chilled water.

The second source of chilled water is from a 100-ton Carrier 30GT air-cooled reciprocating chiller, installed outside the electrical room.

Both Chiller 1 and Chiller 4 ran continuously during the 3-day monitoring period.

The chilled water pumps serving each operating chiller are constant-speed. Modulation of the amount of cooling provided by each cooling coil is achieved by 3-way valves.

The main chilled water plant has two 15-hp chilled water pumps. Only one of them (Pump 2) operated during the 3-day study; it ran continuously. The chilled water pump motors are constant speed.

There are two 25-hp condenser water pumps. Only one of them (Pump 2) ran during the 3-day study. The pump motors are equipped with variable-speed drives.

Figure 9. Chiller 1



Figure 10. Chiller 4

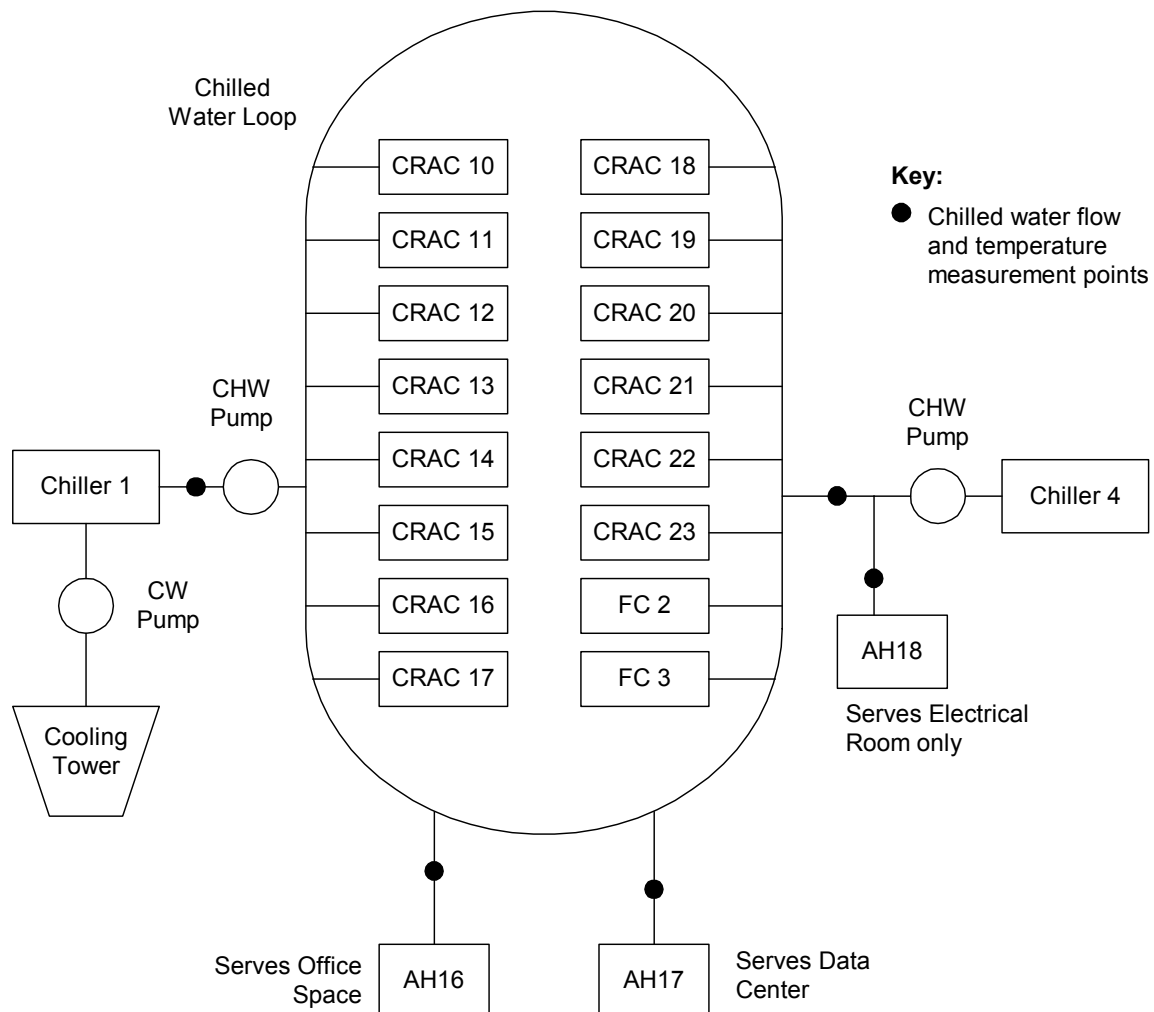


The main plant has a pair of cooling towers. Each cooling tower has one variable-speed fan motor. The condenser water circulates through both towers simultaneously whenever the condenser pump operates, regardless if the fans are running or not. The fan in Cooling Tower 1 ran continuously during the 3-day study; Cooling Tower 2 did not. The main plant is not equipped to do water-side economizing.

The Chiller 4 plant is air cooled, so it does not have condenser pumps or cooling towers. It has two 10 hp chilled water pumps, also numbered 1 and 2. Only Pump 2 operated during the 3-day study. It ran continuously. The pump motors are constant-speed.

The entire HVAC system is controlled and monitored by a Barber Coleman Network 8000 system.

Figure 11. Simplified Schematic of Chilled Water System



Electric power consumption, chilled water flow, and chilled water temperatures⁶ were measured in both chiller plants over two days. Minute-by-minute chiller cooling loads were obtained by multiplying the chilled water flow rate by the chilled water supply and return temperature differential, and by the appropriate conversion factors. The average cooling load is the average of all the 1-minute loads. See Table 6. The outdoor air temperature varied between 40°F and 60°F during this period, and the outdoor relative humidity varied between 60% and 100%⁷. Appendix A contains graphs of the recorded data.

Table 6. Chilled Water Plant Electrical and Load Measurements

Equipment	Nominal Size	Spot / Monitored	Date(s)	Average Value
<i>Chiller 1</i>				
Chiller 1 Power Consumption	n/a	Monitored	1/7/03 – 1/8/03	42.0 kW
Chiller 1 Cooling Load	150 tons	Monitored	1/7/03 – 1/8/03	63.9 tons
Chilled Water Pump 2	15 hp	Spot	1/7/03	8.5 kW
Condenser Water Pump 2	25 hp	Monitored	1/8/03	3.8 kW
Cooling Tower 1 Fan	n/a	Monitored	1/8/03	0.3 kW
<i>Chiller 4</i>				
Chiller 4 Power Consumption	n/a	Monitored	1/7/03 – 1/8/03	36.4 kW
Chiller 4 Cooling Load, Total	100 tons	Monitored	1/7/03 – 1/8/03	47.7 tons
Chilled Water Pump 2	10 hp	Spot	1/8/03	6.7 kW

During the monitored period Chiller 1 drew between 38.9 and 44.5 kW, with an average 42.0 kW. Chiller 1 delivered between 53.5 and 70.1 tons of cooling, with an average of 63.9 tons.

A spot measurement of Chilled Water Pump 2 yielded 8.5 kW.

⁶ These were measured using a Summit Technology PowerSight PS3000 for electric loads, a Controlotron 1010 ultrasonic flow meter for chilled water flow, and a Pace Scientific XR440 Pocket Logger equipped with thermistors to measure the chilled water supply and return temperatures. 1-minute samples.

⁷ Measured with a Pace Scientific XR440 Pocket Logger equipped with a TRH-100 sensor. 1-minute samples. Instrument installed outside building, in shade, near Chiller 4.

Condenser Water Pump 2 is outfitted with a variable speed drive. A variable speed condenser water pump is a good energy saving strategy if the pump serves multiple chillers, or serves a single chiller that has a varying load. Care must be taken to not let the condenser water flow drop below the chiller manufacturer's recommended minimum. The power consumption of Pump 2 was monitored for 3.5 hours on 1/8/03; it averaged 3.8 kW. The data shows a pronounced 1-hour cycle, between approximately 3 kW and 4.5 kW. A chart is included in Appendix A. Facility staff indicated that the variable speed drive uses the pressure drop across the chiller condenser as a control signal, but did not have an immediate explanation for the observed cycling. It appears the control scheme would benefit from a re-examination.

The fan in Cooling Tower 1 drew a relatively constant 0.3 kW during 3 hours of monitoring.

The power draw of Chiller 4 cycled up and down in a regular pattern as it operated. This is typical for reciprocating chillers, as the compressors stage on and off. The minimum value seen during the period of measurement was 26.4 kW, the maximum was 48.4 kW, and the average was 36.4 kW.

A spot measurement of Pump 2 on the Chiller 4 plant yielded 6.7 kW.

AIR HANDLERS

The cooling coils in air handlers AH16, AH17, and AH18 all use chilled water from the chilled water distribution loop. AH16 serves office space which is not part of the data center. AH17 serves only the data center. AH18 serves only the electrical room.

AH16 and AH17 both draw air from a common air-side economizer.

AH17 and AH18 are equipped with Bell & Gossett circuit setters. These devices are basically manually adjustable valves that have an indicator wheel to show valve position, and two access ports for measuring the pressure drop across the valve. Given the valve size, indicator position, and pressure reading, the flow rate can be found in the Bell & Gossett tables.

AH16 is outfitted with a non-adjustable orifice plate that serves the same flow-limiting purpose as the circuit setters. It is equipped with access ports for reading pressure, but has no other identifying information on it. No specification is found in the available mechanical drawings.

The tons of cooling provided by a given cooling coil are obtained by multiplying the chilled water flow rate by the chilled water supply and return temperature differential. Table 7 shows the results.

Table 7. Makeup Air Handler Measurements

Handler	Parameter	Spot / Monitored	Date(s)	Average Value
AH16	Power Consumption	not measured	n/a	n/a
	Cooling Coil Load	Monitored	1/6/03 – 1/8/03	0.0 tons
AH17	Power Consumption	Spot	1/7/03	0.57 kW
	Cooling Coil Load	Monitored CHW Temps; Spot CHW Flow	1/6/03 – 1/8/03	0.0 tons
	Cooling Provided to Data Center	Spot	1/8/03	1.2 tons
AH18	Power Consumption	Spot	1/7/03	4.98 kW
	Cooling Coil Load	Spot	1/6/03	37.5 tons

The power consumption of AH16 was not measured because the office space is not part of this study. The tonnage of AH16 is required so it can be subtracted from the total tons to get the CRAC tons. The measured pressure drop across AH16's orifice plate was very small for the entire monitored period. A chart of the pressure readings is included in Appendix A. For the purposes of this study, the cooling coil load of AH16 is treated as zero. Note that the office space still received cooling via the air-side economizer.

A spot measurement of the pressure drop across the circuit setter⁸ on AH17 yielded a result of 15 gpm of chilled water flow. However, the difference in the monitored supply and return chilled water temperatures showed essentially a zero temperature difference. A chart is included in Appendix A. Like AH16, AH17 used outside air to provide cooling during the monitored period. The cooling delivered to the data center was calculated from the air flow and temperature readings at the supply air diffusers in the data center. See Appendix C for details.

A spot measurement of the pressure drop across the circuit setter on AH18 yielded a result of 78 gpm of chilled water flow. A spot measurement of the chilled water supply and return temperatures gave a differential of 11.55°F. The resulting cooling coil load is 37.5 tons.

⁸ Spot pressure measurements were made with an Alnor HM650 hydronic manometer.

CRAC UNITS

The data center is cooled primarily by twelve floor-mounted Computer Room Air Conditioning (CRAC) units. In addition, there are four fan coil units mounted in the interstitial space above the ceiling tiles. See Figure 2 and Table 1.

All sixteen units use chilled water, have constant-speed fans, and serve only the data center. All units except FC-2 and FC-3 are equipped with electric humidifying units. Electric reheat is used during dehumidification. All sixteen units were running during the period of measurement.

The outside dry bulb air temperature varied between 40 °F and 60 °F during this time, and the outside air relative humidity varied between 60% and 100%. A chart of outside air conditions is included in Appendix A.

Table 8. CRAC Unit Settings

Label	Temperature (°F)		Relative Humidity (%)	
	Setpoint	Tolerance	Setpoint	Tolerance
CRAC-10	70	3	45	3
CRAC-11	70	3	40	3
CRAC-14	70	3	40	3
CRAC-15	70	3	45	3
CRAC-16	70	3	50	3
CRAC-17	70	3	50	3
CRAC-18	70	10	40	15
CRAC-19	70	10	40	15
CRAC-13	70	3	50	3
CRAC-20	70	3	45	3
CRAC-21	70	3	45	3
CRAC-22	69	3	48	3
CRAC-12	70	3	45	5
CRAC-23	70	2	50	5
FC-2	n/a	n/a	n/a	n/a
FC-3	n/a	n/a	n/a	n/a

The measurement team did not measure temperature and humidity in the data center with independent instruments. Ambient conditions as reported by the digital displays on the CRAC units are shown in Table 10.

Figure 12. CRAC 14



The actual cooling tons delivered by each individual CRAC unit was not measured.

Spot measurements were made of the power consumption of all the CRAC units; see Table 9.

Table 9. CRAC Power Consumption Measurements⁹

Label	Nominal Size (tons)	Date & Time	Power Consumption (kW)	Power Factor
CRAC-10	20	1/7/03 12:28	13.72	0.87
CRAC-11	20	1/7/03 10:47	4.81	--
CRAC-14	20	1/7/03 10:56	32.30	0.99
CRAC-15	25	1/7/03 10:59	14.35	0.96
CRAC-16	25	1/7/03 12:25	4.76	0.72
CRAC-17	25	1/7/03 12:14	4.54	0.72
CRAC-18	5	1/8/03 17:30	1.23	0.89
CRAC-19	5	1/8/03 17:37	1.25	0.99
CRAC-13	20	1/7/03 12:22	4.66	0.74
CRAC-20	5	1/7/03 12:17	11.86	0.93
CRAC-21	5	1/7/03 11:03	9.40	0.86
CRAC-22	20	1/7/03 11:10	10.26	0.86
CRAC-12	10	1/7/03 10:52	18.71	0.99
CRAC-23	10	1/7/03 12:20	4.88	0.91
FC-2 & FC-3	?	1/8/03 17:49	1.68	0.61
Total	>215	--	138.4	--

FC-2 and FC-3 are on the same circuit. The units are staged; FC-2 comes on first, and FC-3 comes on only if needed.

The total electric power consumption of the CRAC units was measured as 138.4 kW. This is essentially the sum of the electric power draw of the fan motor in each unit, plus any humidifier or reheat power. The power levels vary significantly from one unit to the other. Table 10 points to an explanation of why this is so.

⁹ All spot measurements were made with a Powersight model PS3000 and clamp-on, 1000-amp current transducers.

Table 10. CRAC Display Panel Readings¹⁰

Label	CW Valve Position (%)	Ambient Temp. (°F)	Ambient Relative Humidity (%)	Heating (%)	De-humid. (%)	Humid. (%)	Time
CRAC-10	0	68	48	33	38	0	10:11
CRAC-11	0	69	48	33	100	0	10:13
CRAC-14	100	66	42	66	100	--	10:16
CRAC-15	23	70	39	0	--	100	10:15
CRAC-16	42	70	46	0	off	--	10:06
CRAC-17	34	70	47	0	off	--	10:07
CRAC-18	n/a	n/a	n/a	n/a	n/a	n/a	n/a
CRAC-19	n/a	n/a	n/a	n/a	n/a	n/a	n/a
CRAC-13	38	69	47	0	off	off	10:20
CRAC-20	0	67	43	50	off	--	11:34
CRAC-21	0	69	44	0	off	--	11:38
CRAC-22	44	69	39	0	--	100	11:37
CRAC-12	100	69	46	66	100	--	10:49
CRAC-23	7	69	45	0	--	100	10:50
FC-2	n/a	n/a	n/a	n/a	n/a	n/a	n/a
FC-3	n/a	n/a	n/a	n/a	n/a	n/a	n/a

In the Main Computer Room, CRAC units 11 and 14 are in 100% dehumidifying mode, and CRAC 15 is in 100% humidifying mode. Likewise, in the Tape Room, CRAC units 12 and 23 are working at cross-purposes. Table 8 offers additional evidence. In the Main Computer Room, the relative humidity tolerance for each CRAC unit is a tight +/- 3%, but CRAC units 11 and 14 have a setpoint of 40% RH, and CRAC units 16 and 17 are set for 50% RH. As would be expected, CRAC units 11 and 14 are in dehumidifying mode.

In general, it appears that the units in dehumidifying mode draw about 2 kW per nameplate ton, those in humidifying mode draw about 0.5 kW per nameplate ton, and those that are neither humidifying or dehumidifying draw about 0.2 kW per nameplate ton. CRAC units 11, 20, and 21 don't follow this pattern, but that may be due to the fact that the power levels and the digital displays were read at different times. The CRAC unit humidifying and dehumidifying functions are likely cycling on and off.

This situation bears further investigation. If the CRAC units are indeed fighting each other over humidity control, rectifying the problem will save significant energy.

¹⁰ All spot measurements were made with a Powersight model PS3000 and clamp-on, 1000-amp current transducers. All readings taken on 1/8/03.

Finally, the investigation team performed airflow efficiency measurements on CRAC units 13 and 16. These units were selected because neither of them was in humidifying or dehumidifying mode. The power consumption for each is essentially due to the fan motor.

Table 11. CRAC 16 Air Flow Efficiency

Grid	Date & Time	Average Velocity	Average Flow Rate	Notes
		fpm	cfm	
1	1/8/03 12:00	598	1860	17" x 30" Filter (16" x 28" Effective Area)
2	1/8/03 12:00	585	1818	17" x 30" Filter (16" x 28" Effective Area)
3	1/8/03 12:00	522	1625	17" x 30" Filter (16" x 28" Effective Area)
4	1/8/03 12:00	598	1861	17" x 30" Filter (16" x 28" Effective Area)
5	1/8/03 12:00	612	1903	17" x 30" Filter (16" x 28" Effective Area)
6	1/8/03 12:00	404	530	8" x 28" Filter (7" x 27" Effective Area)
Total cfm			9597	
kW			4.76	From Table 9
cfm/kW			2020	

Table 12. CRAC 13 Air Flow Efficiency

Grid	Date & Time	Average Velocity	Average Flow Rate	Notes
		fpm	cfm	
1	1/8/03 13:00	695	2161	17" x 30" Filter (16" x 28" Effective Area)
2	1/8/03 13:00	660	2052	17" x 30" Filter (16" x 28" Effective Area)
3	1/8/03 13:00	594	1849	17" x 30" Filter (16" x 28" Effective Area)
4	1/8/03 13:00	689	2144	17" x 30" Filter (16" x 28" Effective Area)
Total cfm			8206	
kW			4.66	From Table 9
cfm/kW			1762	

LIGHTING

Lighting in the data center consists of T-8 fluorescent lamps with electronic ballasts. There are four lamps and two ballasts per fixture. There are no occupancy sensors. According to facility staff, the lights are on continuously. The total lighting power was estimated by counting the number of fixtures and multiplying by 32 Watts per lit lamp and 8 Watts per ballast.

Table 13. Lighting Power

	Floor Area (Sq. Ft.)	Fixture Count	No. of Lamps Not Lit	kW
Main Computer Room	6,450	87	16	12
Network Room	900	16	0	2
New Print Room	3,680	26	4	4
Old Print Room	2,260	44	3	6
Tape Room	1,580	14	0	2
Phone Room	450	n/a	n/a	n/a
Totals:	15,320	187	23	26

The lighting power density is 1.7 Watts per square foot of gross floor area. This is much higher than the typical 1.0 W/sf for office environments.

SUMMARY MEASUREMENTS AND METRICS

SUMMARY MEASUREMENTS

The data center net cooling load is determined indirectly. All of the cooling provided by Chiller 1 serves the data center. The cooling coil load of air handler AH18 is subtracted from the total tons provided by Chiller 4; the remainder serves the data center. Finally, the cooling provided by makeup air handler AH17 is included. Table 14 shows the result.


Table 14. Total Data Center Cooling Load

Cooling Provided By	Average Value
Chiller 1	63.9 tons
Chiller 4 (not including AH18 load)	10.2 tons
AH17	1.2 tons
Total	75.3 tons

Table 15 brings together all the electrical measurements for the data center.


Table 15. Summary of Electrical Measurements

Item	Remarks	Value	Percent
Computer Equipment Loads	Total UPS output. (All data center equipment is served by PDUs; all PDUs are served by UPS system.)	190.0 kW	41%
UPS Losses	Total UPS input minus total UPS output.	30.6 kW	7%
HVAC – Air Movement & Humidity Control	Electrical energy for all CRAC and Fan Coil Units in data center, plus AH17 fan energy.	139.0 kW	31%
HVAC – Chiller & Pump Power	Power consumption of chillers, pumps, and cooling tower fan. Chiller 4 power prorated to account for AH18 load.	69.1 kW	15%
Lighting	Calculated from fixture count.	26.2 kW	6%
Total Energy Use	--	454.9 kW	100%

Figure 13. Data Center Energy Balance

The computer equipment load served by the UPS system is 41% of the data center energy usage. The energy used to cool the data center outweighs this; pumping and chiller energy accounts for 15% of the total, and air movement and humidity control amounts to 31%, for a total of 46%. See Figure 14 for a comparison to the other facilities studied as part of the Data Center Load Characterization Project. UPS losses account for 7% of the data center energy consumption, and the lights account for the remaining 6%.

Figure 14. HVAC Power Comparison



ELECTRICAL CONSUMPTION METRICS

Table 16 addresses the issue of data center load density. The most commonly used metric among mission critical facilities is the computer load density in Watts consumed per square foot (W/sf). However, even in a prototypical data center filled entirely with closely spaced racks of similar equipment, the choice of what to use as square footage is not always consistent between analysts, and can be a source of confusion.¹¹ In the case of Data Center Facility 9, the choice is further complicated by the fact that the various rooms that make up the data center have various kinds of equipment in them. For the purposes of this report, no breakdown of the different types of computer equipment is performed. All the types are treated alike, and termed “computer equipment”.

¹¹ See “Data Center Power Requirements: Measurements from Silicon Valley”, by Mitchell-Jackson, Koomey, Nordman, & Blazek, December 2001. It is available on the web at http://enduse.lbl.gov/Info/Data_Center_Journal_Article2.pdf.

Table 16. Electrical Consumption Metrics

Metric		Value	Units
Data Center Gross Area		15,300	sf
Computer Equipment Footprint		1,970	sf
Computer Load Density based on Data Center Gross Area		12.4	W/sf
Computer Load Density based on Computer Equipment Footprint		96.4	W/sf
Current Occupancy	Main Computer Room	40	%
	Network Room	70	%
	New Print Room	30	%
	Old Print Room	80	%
	Tape Room	100	%
	Phone Room	100	%
Computer Load Density based on Occupied Floor Area		23.3	W/sf
Computer Load at 100% Occupancy		428	kW


“Data Center Gross Area” is the entire floor area of the data center. Per the Uptime Institute Definitions, this gross floor space is what is typically used by facility engineers in calculating a computer load density (W/sf).¹²

“Computer Equipment Footprint” is the portion of the floor area immediately underneath the equipment. For computer load, we use the total UPS output of 190.0 kW.

The computer load density based on the gross area is 12.4 W/sf. This is well below current typical data center densities, which are in the range of 30 to 50 W/sf. For comparison with other facilities measured as part of the Data Center Load Characterization Project, see Figure 15.

¹² Users look at watts per square foot in a different way. With an entire room full of communication and computer equipment, they are not so much concerned with the power density associated with a specific footprint or floor tile, but with larger areas and perhaps even the entire room. Facilities engineers typically take the actual UPS power output consumed by computer hardware and communication equipment in the room being studied (but not including air handlers, lights, etc.) and divide it by the gross floor space in the room. The gross space of a room will typically include a lot of areas not consuming UPS power such as access aisles, white areas where no computer equipment is installed yet, and space for site infrastructure equipment like Power Distribution Units (PDU) and air handlers. The resulting gross watts per square foot (watt/ft²-gross) or gross watts per square meter (watt/m²-gross) will be significantly lower than the watts per footprint measured by a hardware manufacturer in a laboratory setting.

Figure 15. Computer Load Densities



The computer load density based on the computer equipment footprint is 96.4 W/sf.

Occupancy was determined by examining the electronic drawings for the data center. The computer equipment in the main computer room is clustered most closely together in the north end of the room. If the main computer room was built out such that the entire room achieved this spatial density, there would approximately 2.5 times as much equipment in the room. Thus the current occupancy is 40%.

A similar calculation for the Network Room, New Print Room and the Old Print Room yields 70%, 30%, and 80%, respectively.

The Tape Room and Phone Room are fully occupied with their respective types of equipment.

Using the current average computer load density for the entire data center and extrapolating to full occupancy yields 428 kW. This is within the 500 kVA capacity of each UPS.

HVAC EFFICIENCY METRICS

Since the packing of data centers and computer types are site specific, a more useful metric for evaluating how efficiently the data center is cooled can be represented as a ratio of cooling power to computer power.

Another metric is the “theoretical cooling load”. It is the sum of the computer, lighting, and CRAC electrical loads. All of these loads equate to heat that must be removed from the room. (Though there is a small amount of human activity, this is insignificant compared to the other loads.)

Chiller efficiency is usually presented as the ratio of chiller power at full load to the tons of cooling provided at full load, in units of kW/ton. Chilled water plant efficiency is similar, but it includes the power consumption of the cooling tower and pumps as well. HVAC system efficiency adds the power consumption of the air handlers and CRAC units.

Table 17. HVAC Efficiency Metrics

Metric	Average Value	Units
Cooling kW / Computer Load kW	1.10	--
Theoretical Cooling Load	101	tons
Cooling Provided by HVAC System	75.3	tons
Chiller 1 Efficiency	0.66	kW/ton
Chiller 4 Efficiency	0.76	kW/ton
Chilled Water Plant Efficiency	0.93	kW/ton
HVAC System Efficiency	2.81	kW/ton

The data center uses 10% more energy to provide cooling for the computers, than the computers themselves consume. This figure does not include the efficiency loss in the UPS, or the energy required to keep the UPS cool.

The theoretical load is 34% higher than the measured cooling delivered by the HVAC system. It is likely that some of heat generated in the data center is escaping through the roof and in to adjoining spaces in the facility.

The efficiency of Chiller 1 averaged 0.66 kW/ton. This is typical for an older, lightly loaded centrifugal chiller, but newer variable-speed chillers can be much more efficient. Efficiencies of 0.3 kW/ton or less are possible.

The efficiency of Chiller 4 averaged 0.76 kW/ton. This is less efficient than Chiller 1, but it is fairly good performance for an air-cooled reciprocating unit.

The chilled water plant as a whole – both chillers and their associated chilled water pumps, the condenser water pump on Chiller 1, and the cooling tower – is less efficient than either chiller by itself, but not by a huge amount. The variable speed drives on the condenser pump and the cooling tower keep the power consumption of those devices relatively low.

The overall HVAC system efficiency is 2.81 kW/ton. As shown previously in Figure 13, two-thirds of the HVAC power consumption is due to the CRAC units. This portion of the energy use can probably be reduced by addressing the humidifying/dehumidifying control issue in the data center.


V. Energy Efficiency Recommendations

GENERAL GUIDELINES FOR FUTURE CONSTRUCTION

Efficient Chilled Water System

Water cooled chillers offer enormous energy savings over air cooled chillers, particularly in dry climates, because they take advantage of evaporative cooling. Since the chiller is being cooled by lower temperature media, it can reject heat more easily, and does not have to work as hard. Though the addition of a cooling tower adds maintenance costs associated with the water treatment, we have found that the energy savings outweigh the maintenance costs. Within the options of water cooled chillers, variable speed centrifugal are the most energy efficient, because they can operate very efficiently at low loads. The graph below compares the energy performance of various chiller types.

Comparison of Typical Chiller Efficiencies over Load Range

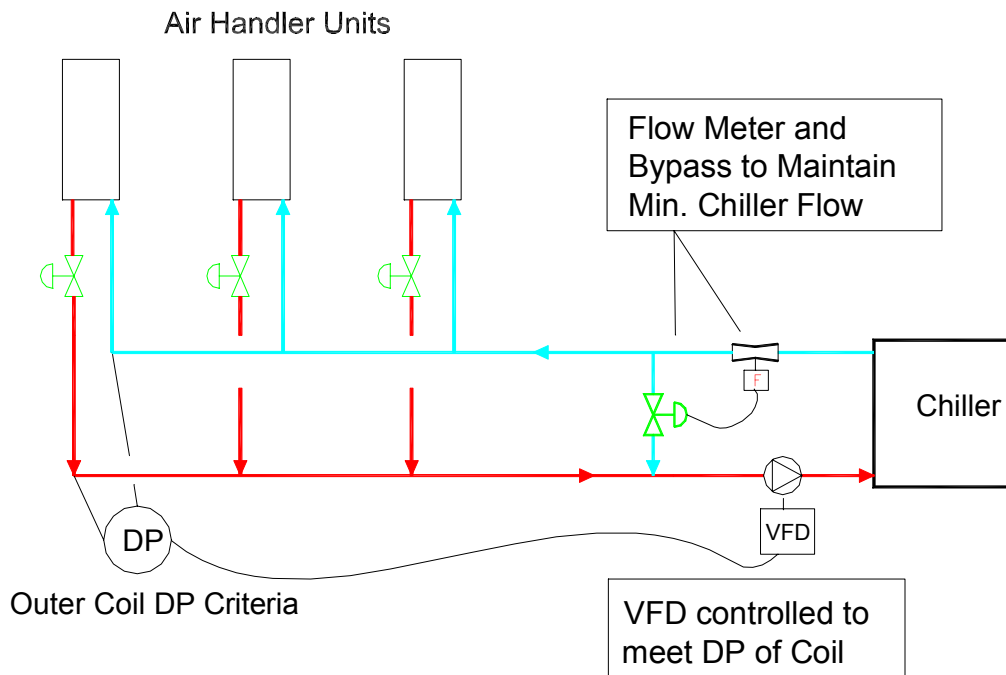


Chiller 1	250-Ton, Screw, Standard Efficiency, Air Cooled
Chiller 2	216 Ton, Screw, Water Cooled
Chiller 3	227-Ton, Centrifugal, Constant Speed, Water Cooled
Chiller 4	227-Ton, Centrifugal, Variable Speed, Water Cooled

Though there are efficient air cooled chillers, the larger size of water cooled chillers has resulted in more care given to efficiency and life cycle costs compared to air cooled chillers.

The selection of the auxiliary equipment, including the cooling tower, pumps, and pumping strategy should also be considered carefully. For example, variable speed fans on cooling towers allow for optimized cooling tower control. Premium efficiency motors and high efficiency pumps are recommended, and variable speed pumping is a ripe opportunity for pump savings. Variable pumping strategies can be achieved in a primary/secondary scheme, where the primary pumps operate at constant speed and directly feed water to the chiller, and the secondary pumps are variable speed and serve the air handling units. A more energy efficient scheme is primary-only variable speed pumping strategy. Pumping savings are based on the cube law: pump power is reduced by the cube of the reduction in pump speed, which is directly proportional to the amount of fluid pumped.

A primary only variable pumping strategy must include a bypass valve that ensures minimum flow to the chiller, and the use of two-way valves at the air handling units in order to achieve lower pumping speeds. The control speed of the bypass valve should also meet the chiller manufacturers recommendations of allowable turndown, such that optimum chiller efficiency is achieved.¹³ The diagram below describes the primary-only variable speed pumping strategy.



¹³ This basically means that the flow through the chiller should be varied slow enough such that the chiller is able to reach a quasi-steady state condition and able to perform to its maximum efficiency.

Air Management

The standard practice of cooling data centers employs an underfloor system fed by CRAC units. There are a number of potential problems with such systems: an underfloor system works on the basis of thermal stratification. This means that as the cool air is fed from the underfloor, it absorbs energy from the space, warming up as a result, and rises. In order to take advantage of thermal stratification, the return air must be collected at the ceiling level. CRAC units often have low return air grills, and are therefore, simply recirculating cool or moderately warmed air. Furthermore, they are often located along the perimeter of the building, and not dispersed throughout the floor area, where they can more effectively treat warm air. One alternative is to install transfer grills from the ceiling to the return grill. Another common problem with underfloor supply is that the underfloor becomes congested with cabling, increasing the resistance to air flow. This results in an increase in fan energy use. A generous underfloor depth is essential for effective air distribution (we have seen 3 feet in one facility).

An alternative to underfloor air distribution is high velocity overhead supply, combined with ceiling height return. A central air handling system can be a very efficient air distribution unit. Design considerations include using VFDs on the fans, low pressure drop filters, and coils. An additional advantage of a central air handling system is that it can be specified with an economizer function. With the favorable climate in the Bay Area, economizing can reduce the cooling load for a majority of the hours of the year.

Another common problem identified with CRAC units is that they are often fighting each other in order to maintain a constant humidity setpoint. Not only is a constant humidity setpoint unnecessary for preventing static electricity (the lower limit is more important), but it uses extra energy. A central air handling unit has a better ability to control overall humidity than distributed CRAC units.

Air Management – Rack Configuration

Another factor that influences cooling in data centers is the server rack configuration. It is more logical for the aisles to be arranged such that servers' backs are facing each other, and servers' fronts are facing each other. This way, cool air is drawn in through the front, and hot air blown out the back. The Uptime Institute has published documents describing this method for air management.¹⁴

Commissioning of New Systems and Optimized Control Strategies

Many times the predicted energy savings of new and retrofit projects are not fully realized. Often, this is due to poor and/or incomplete implementation of the energy efficiency recommendations. Commissioning is the process of ensuring that the building systems perform as they were intended to by the design. Effective commissioning actually begins at the design stage, such that the design strategy is critically reviewed. Either the design engineer can serve as the commissioning agent, or a third party commissioning agent can be hired. Commissioning differentiates from standard start-up

¹⁴ <http://www.upsite.com/TUIpages/whitepapers/tuiaaisles.html>

testing in that it ensures systems function well relative to each other. In other words, it employs a systems approach.

Many of the problems identified in building systems are often associated with controls. A good controls scheme begins at the design well. In our experience, an effective controls design includes 1) a detailed points list, with accuracy levels, and sensor types, and 2) a detailed sequence of operations. Both of these components are essential for successfully implementing the recommended high efficiency chilled water system described above. Though commissioning is relatively new to the industry, various organizations have developed standards and guidelines. Such guidelines are available through organizations like the Portland Energy Conservation Inc., at www.peci.org, or ASHRAE, Guideline 1-1996.

Lighting Controls

Lighting controls such as occupancy sensors may be appropriate for areas that are infrequently or irregularly occupied. If 24-hour lighting is desired for security reasons, scarce lighting can be provided at all hours, with additional lighting for occupied periods.

SPECIFIC RECOMMENDATIONS

The chillers studied in this report are due to be removed soon, so no specific recommendations are provided for them.

Revisit Humidity Control Strategy

The CRAC units are programmed with humidity setpoints that vary from 40% to 50% RH, but most of the units are also programmed with relatively tight tolerances (+/- 3% RH). Observations indicate the CRAC units are working against each other as they attempt to meet their respective humidity setpoints. This consumes a significant amount of energy.

Humidity control is an important issue in the print rooms, as paper forms will tear or jam in the printers if the relative humidity is out of range. Humidity may also be an issue with the magnetic data storage tapes. If the remaining computer equipment can tolerate a wider range of humidity, though, relaxing the humidity standard will help the CRAC units stop competing with each other. Additionally, revisit the humidity setpoints. It is probably not necessary to have some CRAC units set for 40% RH and others in the same room set for 50% RH.

Run Fewer CRAC Units

The nominal cooling capacity of all the CRAC units combined is more than 215 tons. Yet, during the monitoring period, the data center called for an average of only 74 tons. All the CRAC units ran constantly during the monitored period, and all of them have constant-speed fans. If it is possible to turn off some CRAC units without creating “hot

spots”, then it will save the fan energy. The CRAC units that continue to operate will increase their use of chilled water to meet the total cooling demand .

Rearrange Perforated Floor Tiles for More Effective Cooling

Some of the floor tiles in the data center are perforated, to allow the cooling air to rise from the space under the floor. The investigation team noticed that many of the perforated tiles could be rearranged to more efficiently cool the computer equipment. In particular, perforated tiles should be placed in front of server racks, not behind. The cooling fans inside the servers typically draw air from the front of the rack and eject it out the back. Directing the cooling air that is coming from the floor to rise in front of the rack will provide the optimum cooling effect.

Consider Using Outside Air to Cool the Data Center

Air handler AH17 provides ventilation air to the data center through three ceiling diffusers. This air handler is capable of using outside air when the outside air is cool enough, thereby avoiding the need for the chilled water plant to provide chilled water for cooling. This concept can be employed for active cooling, not just ventilation. The outside air in Sacramento is cool and dry enough through much of the year, to allow using it directly for cooling. This would require a new air handler and ductwork to serve the data center, so a cost-benefit analysis would be called for as a first step.

Examine UPS Management Strategies

Each UPS system is large enough to carry the data center load by itself. Running one UPS at a higher load factor will yield better efficiency. However, the battery bank in each system will not be maintained properly if the respective system is not active. Perhaps it would be possible to alternate the two UPS systems in a lead/lag arrangement, one on and one off, switching them frequently enough that the batteries are not negatively impacted?

Reduce Lighting

The data center is 15,300 square feet, and uses a constant 26 kW of lighting energy. If electricity costs \$0.10/kWh, this represents some \$23,000 per year. The lighting energy density works out to 1.7 W/sf, which is much higher than the 1.0 W/sf found in a typical office environment. In addition, some of the rooms are occupied intermittently.


Consider the following measures:

- a) Install occupancy sensors to turn off most or all of the lights when the rooms are unoccupied.
- b) Add task lighting in appropriate areas, and disable a portion of the overhead lights.


Appendix A

Charts of Measured Data


Facility 9 Data Center Outside Air Conditons




Facility 9 Data Center Chiller 1 Temperatures




Facility 9 Data Center Chiller 1 Power




Facility 9 Data Center Chiller 1 Flow & Tonnage




Facility 9 Data Center Chiller 1 Efficiency




Facility 9 Data Center Cooling Tower Fan Power




Facility 9 Data Center Cooling Tower Pump Power




Facility 9 Data Center Chiller 4 Temperatures




Facility 9 Data Center Chiller 4 Power




Facility 9 Data Center
Chiller 4 Flow & Tonnage, Downstream of AH18




Facility 9 Data Center Chiller 4 Efficiency




Facility 9 Data Center
AHU 16 Cooling Coil Water Temperatures



Facility 9 Data Center
AHU 16 Cooling Coil Flow Setter Pressures



Facility 9 Data Center
AHU 17 Cooling Coil Water Temperatures



Appendix B

Measurement of UPS Efficiency

UPS Measurements

ID	Input or Output	Date & Time	Source	Measurements								Calculations			
				Volts			Amps			kW	pf	Power	Effic.	Load Factor	Loss
				A-B	B-C	C-A	A-B	B-C	C-A			kW	%	%	kW
UPS-1	Input	1/7/03 14:12	UPS Display	489	486	487	267.2	259.0	289.0	-	-	227.1	89.6	40.7	23.6
	Input	1/7/03 13:53	Powersight	-	-	-	242.4	229.5	244.7	196.4	0.99				
	Output	1/7/03 14:12	UPS Display	-	-	-	-	-	-	203.5	-				
UPS-2	Input	1/7/03 14:10	UPS Display	483	482	482	164.4	164.8	164.8	-	-	35.8	47.3	3.4	18.8
	Input	1/7/03 14:02	Powersight	480	490	499	166.0	169.0	164.4	36.0	0.26				
	Output	1/7/03 14:10	UPS Display	-	-	-	-	-	-	17.0	-				
Combined	Input											262.8	83.9	16.8	42.3
	Output											220.5			

“Powersight” refers to a Powersight PS3000 power meter with 3000-amp flexible current transducers. The UPS output power wires were not accessible for direct measurement.

The above table compares Powersight readings to the UPS digital displays. The UPS displays show input power in terms of volts and amps, not kW. The input power is calculated as the average of volts times amps on each leg, times the square root of 3, times the power factor as measured by the Powersight.

The UPS display shows output power directly in kW.

There is close agreement between the Powersight and the digital display on UPS 2.

There is a 16% disagreement between the Powersight and the display on UPS 1. This is probably due to the time difference between readings; it is possible that additional computer equipment came on line between the time of the Powersight reading and the time of the display reading. The display value of 227.1 input kW is used to calculate efficiency, load factor, and loss.

Appendix C

Cooling Provided by Air Handler AH17

Air Handler AH17 Measurements

Supply Air

Diffuser/Grill Location	Date & Time	Velocity Measurements (fpm)					Flow	Temperature
		1	2	3	4	Ave	cfm	deg F
Tape Room (1 of 2)	1/8/03 13:30	82	91	83	-	85	341	56
Tape Room (2 of 2)	1/8/03 13:30	84	82	82	-	83	331	57
Main Computer Room (1 of 1)	1/8/03 13:30	82	78	82	77	80	319	55
Total:							991	
Ave:								56

Return Air

Main Computer Room (1 of 4)	1/8/03 13:40							71
Main Computer Room (2 of 4)	1/8/03 13:40							70
Main Computer Room (3 of 4)	1/8/03 13:40							69
Main Computer Room (4 of 4)	1/8/03 13:40							70
Ave:								70

Tons: 1.2

All of the air from air handler AH17 is delivered to three ceiling diffusers in the data center. Figure 2 in the report shows their locations.

Air velocity was measured with a Shortridge ADM-860 Airdata multimeter.

Temperatures were measured with a Raytek MiniTemp handheld infrared thermometer.

Appendix XI. “NY Data Center Energy Benchmarking and Case Study, Facility 10”; NYSERDA Case Study, Syska & Hennessy; LBNL – William Tschudi, 2003



**ERNEST ORLANDO LAWRENCE
BERKELEY NATIONAL LABORATORY**



**SYSKA HENNESSY
GROUP**

NY DATA CENTER ENERGY BENCHMARKING AND CASE STUDY

MARCH, 2003

PRINCIPAL INVESTIGATOR: WILLIAM TSCHUDI - LBNL

SUBCONTRACTORS: PAUL LIESMAN - SYSKA AND HENNESSY

JAMES McENTEGGART – SYSKA AND HENNESSY

CONTENTS

EXECUTIVE SUMMARY	3
INTRODUCTION.....	6
SITE OVERVIEW	6
ENERGY USE	7
ANALYSIS	14
APENDICES	26
Appendix A - Fluid Thermal Energy Meter Data	27
Appendix B - Power Monitoring Report.....	31
Appendix C - Data Center References	48

I. Executive Summary

The New York State Energy Research and Development Authority (NYSERDA) sponsored this project to study energy use in a New York data center. The study focuses on energy efficiency issues in the selected data center, determines energy end use, and looks for energy efficiency opportunities.

A goal of this project was to provide the actual energy intensity in an operating data center in New York. In addition a particular focus of the study was to study the control systems for various energy intensive facility systems.

This information together with other studies will provide insight into the distribution of electrical power with the data center and the overall electrical demand for data center facilities. Energy benchmark data for a sufficient number of data centers will also eventually help to identify current best practices, and determine efficiency and reliability improvement areas. Additional case studies benchmarking energy use in California data centers were performed. These projects are developing a more robust set of benchmarks and efficiency recommendations.

The facility selected for the study was a mixed-use facility, with several data center areas distributed throughout two buildings. Two large main buildings were in excess of 415,000 sq. ft. and also contained a large amount of office space, a cafeteria, and other support spaces. Chilled water was provided by a decentralized chilled water plant consisting of five chillers of varying sizes and ages, tied to a centralized condenser water system. Chilled water was pumped into a common header for distribution throughout the entire facility. Base load chilled water requirements were met using a relatively new 1,000 ton chiller. The new chiller was assumed to be the most efficient (the lowest kW/ton of chilled water), but this was not verified by direct measurement during this study. Due to the size and complexity of the heating, chilled water, and other building infrastructure systems serving the entire complex, they were beyond the scope of this study. The main interest focused on electrical power consumption within the defined data center area.

The data centers within the selected facility contained a large variety of computing equipment for various uses. The facility is mainly used for data recovery purposes involving multiple customers with a wide range of computing equipment. As such, the operations may not have been as uniform as some other types of data centers, such as web hosting or dedicated data processing. Through investigation and from power use profiles collected in our measurements, the data recovery activity was observed to cause variation in power usage not typically encountered in other data centers.

Figure 1 is a representative time plot showing some variation in electrical load. The site team did not attempt to determine the cause of variations in computing loads.

Energy metrics were developed to determine and track energy intensity for the overall facility and for the computing space. These metrics allow comparison to other datacenters and provide indicators of the performance of individual systems and components.

The benchmarks obtained in this study are useful to the host facility for several purposes:

- ◆ Providing a baseline to track performance over time
- ◆ Identifying the most energy intensive systems and components
- ◆ Uncovering operating and maintenance problems
- ◆ Finding energy and reliability operating and retrofit improvement opportunities
- ◆ Comparing performance to benchmarks observed elsewhere
- ◆ Determining energy intensity trends in computing equipment over time
- ◆ Establishing efficiency improvement goals based upon benchmark information
- ◆ Establishing operating and design targets for future projects

As more robust benchmark data is available (a statistically significant set of data), it is expected that best practices will emerge. In addition, the host site, being part of a large national firm, should be able to compare its performance to its other data center facilities within the firm as well as to the data centers benchmarked by LBNL or others. This may provide awareness of opportunities for continual improvement.

The host site proved to be a difficult case study for several reasons. The facility, being older and modified by years of additions and renovations provided a challenge to isolating areas/systems of interest. A significant amount of office space and other support spaces made isolation of just the data center systems and space a difficult task. In addition, information that is often provided in building information and management systems in more modern facilities was not readily available. The mixed use of the facility, combined with complex systems (such as the chilled water system) that were added to over the life of the building, necessitated taking a simplified approach to




Figure 1 – Load variation

evaluating systems. To obtain useful data at a more detailed level, the measurement team defined a smaller “control volume” to study the energy end-use in one representative computing area. The study was further constrained by the fact that the data collection team had difficulty obtaining site information due to staff reassignment at the facility. Nonetheless, useful data and observations were eventually obtained.

The primary area of interest was the HVAC system serving the data center areas, but the study also included other data, such as other typical office loads where the data was available or easily obtainable. Whole facility energy use was obtained and then systems serving the smaller control volume were evaluated to as detailed a level as practical. Unfortunately, the original design information was not available due to the age and history of the facility.

The accuracy and completeness of data varies, based upon the measurement methods, access, and ease of measurement. Nonetheless, the data is sufficiently accurate to determine energy intensities and is useful for other observations concerning the facility.

The energy intensity due to the computing equipment in a defined area of the data center (control volume) was calculated based in part upon measured power use along with simplifying assumptions. The intensity values are useful for trending electrical load growth as computing equipment evolves and, in the case of this facility, as computing equipment changes to satisfy customer needs. To quantify the maximum electrical intensity if the center were full (using the current mix of computing equipment) a qualitative assessment of the percent occupied was made (or how full the data center was). A comparison to benchmark data from other data centers in the study is provided in figure 2. In this figure, the load density due to only the computer load is plotted and averaged approximately 25 W/sf. This case study is facility 10 on the graph.




Figure 2 Measured Computer Load Density

A number of observations for possible efficiency improvements or further study for this facility are provided in this report. The observations were not meant to be in response to a comprehensive energy audit, but rather represented specific opportunities for improvement, for further study, or for use in future modifications or new construction.

II. Introduction

In order to gain a better understanding of the energy requirements associated with the increasing use and centralization of data processing equipment in specialized facilities known as “Data Centers” the New York State Electric Research and Development Association (NYSERDA) commissioned Lawrence Berkley National Laboratory (LBNL) to conduct a study of a Data Center in New York to determine the current operating loads and to identify energy efficiency opportunity in the facility. A data center facility volunteered their site in lower New York state for this study. To assist in collecting site data and obtaining measured electrical use for the facility, LBNL contracted with Syska and Hennessy, a data center design firm located in New York City.

III. Site Overview

A multi-national corporation with a data center in Orange County, NY volunteered to participate in the study. This location is a mixed-use facility consisting of office space as well as computing space. One of the functions of the facility is to provide customers with a facility to enable processing and management for disaster recovery. In this regard, the host site or the customer may provide computing equipment. As a result, the mix of computing equipment is diverse and changes frequently. The facility includes two buildings totaling 415,000 square feet (ft²). Of this total 119,000 ft² is classified as technical space (raised floor) whose primary function is the support of data processing equipment. The remaining space is office, cafeteria, and other support space (equipment rooms, supply storage, etc.)

Some general information regarding the function and capacities of this facility are as follows:

- ◆ The two main buildings were constructed in 1972 and 1982 respectively. Additions and renovations occurred over the life of the buildings. A central chilled water central plant provides cooling for the campus. Cooling for the office spaces is supplied from the same system as the data center so that isolating the cooling for the data center spaces only is very difficult.
- ◆ Two Redundant 69kV Electrical Feeds/facility draws approximately 2800 kW
- ◆ Two Independent On-site Substations for Site Service – 10.5 MVA each
- ◆ Three 7.5MW Gas Turbine Back Up Electrical Generators

- ◆ UPS Systems of 1500 KVA and 3000 KVA for Buildings 001 and 002 respectfully
- ◆ Distributed Chiller Plant with 3000 ton capacity –
 - Chiller (electric) capacity of 3400 tons and
 - Cooling tower capacity of 3000 tons.
 - Decentralized chilled water plant consisting of five chillers of varying sizes and ages, tied to a centralized condenser water system.
 - Chilled water pumped into common header for distribution throughout the entire facility.
 - Base load met using newer 1000 ton chiller.
- ◆ Central Boiler Plant with 800 BHP capacity (24,000#/hr) – 2 – 400 BHP Oil fired steam boilers
- ◆ 12,000 gallons of emergency water storage
- ◆ Energy Management Systems – Johnson Controls (METASYS); Westinghouse (INCOM LIGHTING)
- ◆ No natural gas service

IV. ENERGY USE

Historical Data

The host facility staff routinely tracks total facility electric use. This information was provided for the study. As expected, the total facility energy use is relatively constant. Monthly variations in electrical consumption are relatively small and are expected due to the nature of the facility (frequently adding or removing customers and their computing load), and due to weather variations affecting energy use in the larger non-critical areas of the facility (office space and cafeteria). Figure 3 below illustrates the total facility electrical energy use for a one-year period.




Figure 3 Monthly Electricity Use

Although the focus of this study was efficiency, we also observe that the average kW cost varied significantly probably due to demand charges. Figure 4 illustrates the variation in average electricity cost. This suggests that demand reduction strategies such as thermal storage, resetting temperature limits, use of free cooling, etc. may be attractive. Studying these opportunities, however, was beyond the scope of this study.




Figure 4 Average kWh cost

Further analysis of the historical data – calculating the energy intensity associated with the various facility spaces - shows that the data center computing equipment energy intensity overall is approximately 16 W/SF for the whole facility, and has been relatively constant. To calculate this metric, the total electricity serving the computing equipment is divided by the area of raised floor.

Figure 5 illustrates the energy intensities attributable to office space, infrastructure space, and UPS (computer equipment). Infrastructure (facility systems) energy intensity dramatically improved in 1998-1999 and has remained relatively constant since that time, as have the office loads. In this figure, the energy intensity is calculated based upon the respective square footage. Tables 1 and 2 present this detailed information by year.




Figure 5 Facility energy intensity by end use for six years

Another useful breakdown is in the percentage of computing power consumption relative to the other facility loads and the relative square footage of each. This is illustrated in figure 6.




Figure 6 Percentages of UPS, Office and Infrastructure

This illustrates that approximately 53% of the facility electrical use is for computing equipment that occupies 29% of the facility.

Power Consumption (MWH and % of Total)									
	1997			1998			1999		
	MWH	% of Total	W/SF	MWH	% of Total	W/SF	MWH	% of Total	W/SF
Office Loads	16161.43	28%	6.93	8883.93	21%	3.81	6557.99	20%	2.81
Infrastructure Loads	26174.42	46%	100.95	17021.83	40%	65.65	8540.02	27%	32.94
UPS Loads	15127.33	26%	14.51	16249.87	39%	15.59	17076.66	53%	16.38
Total	57463.18	100%	15.81	42155.63	100%	11.60	32174.67	100%	8.85

Table 1 Facility power consumption 1997-1999

Power Consumption (MWH and % of Total)									
	2000			2001			2002(Annualized)		
	MWH	% of Total	W/SF	MWH	% of Total	W/SF	MWH	% of Total	W/SF
Office Loads	6514.45	20%	2.79	6257.81	21%	2.68	5940.97	19%	2.55
Infrastructure Loads	9048.16	28%	34.90	8439.68	28%	32.55	9022.83	28%	34.80
UPS Loads	16442.39	51%	15.77	15163.71	51%	14.55	16960.50	53%	16.27
Total	32005.00	100%	8.80	29861.19	100%	8.21	31924.30	100%	8.78
Building Space Distribution:									
Office Space	266400 Ft ²								
Infrastructure	29600 Ft ²								
Technical Space (UPS Supported)	119000 Ft ²								
Total	415000 Ft ²								

Table 2 Facility power consumption 2000-2002

The overall facility average and peak power density is provided in table 3 below. While these values are not particularly of interest when considering energy use of data center spaces alone, they are useful to the owner of a mixed-use facility such as this to trend overall energy changes due to mix of spaces, energy intensity of new computing equipment, mix of customer equipment, and efficiency of facility systems.

Calculated Facility Peak Electrical Power Density by Year						
	1997	1998	1999	2000	2001	2002
Facility Power Density (W/SF)	15.81	11.60	8.85	8.80	8.21	8.78
Avg Pwr Demand (KW)	6611	4764	3706	3652	2669	2832
Peak Pwr Demand (KW)	7993	6824	4727	6370	4193	4808
<i>Calculated Peak Pwr Density (W/SF)</i>	<i>19.12</i>	<i>16.62</i>	<i>11.29</i>	<i>15.34</i>	<i>12.90</i>	<i>18.10</i>

Table 3 Average and Peak facility energy intensity

Additionally, 2001 monthly energy use data provides insight into the amount of computing load compared to the total load. Figure 6 provides this data. The whole building load (demand) varied less than $\pm 3.3\%$ from the average during this period, confirming the relatively constant computing load. This small variation can be attributed to weather influence and changes in the amount of computing equipment.




Figure 6 Computing Load vs total Load

The trends for energy end use can be seen in figure 7.




Figure 7 Whole building energy end use for six years

Approach for Data Center – Use of Control Volume

In order to study the energy use for the “data center” area, it was necessary to define a control volume that was manageable for this study. A theoretical control volume boundary was established for one area (a portion of the fourth floor of building 1). Energy flow into and out of the area could then be measured and/or approximated.

The area for study was determined after review of the facility plans and discussion with the facility staff. Isolation of a section of the top floor, which consists primarily of technical (data center) space, represented a good control volume to study.

A study of the facility drawings revealed that the upper floor is served by a single chilled water loop for cooling, and three primary electrical feeds. With this information a data collection plan was formulated. The goal was to measure the primary energy inputs into the control volume. A schematic diagram illustrates the chilled water system (figure 8) below:

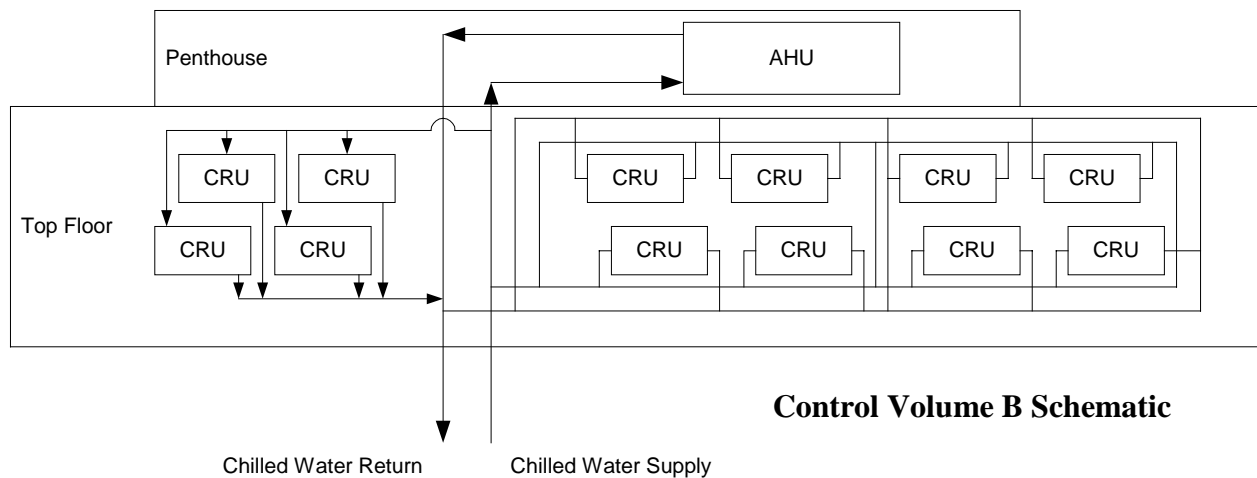


Figure 8 Chilled water system schematic drawing

The data collection effort was coordinated with the staff at the facility and on August 21, 2002 three Dranetz-BMI power analyzers were installed in addition to a Controlotron Energy Meter for measurement of the electrical and chilled water input into the control volume. These devices record the data electronically through the use of an on board computer and memory.

During the installation the team realized that the intended installation location for the Controlotron was not suitable for proper data collection due to lack of a suitable straight run of pipe. Changes in direction in the piping create turbulence, which interferes with the

Figure 9 Chilled Water Piping




operation of the ultrasonic transducers in the flow meter used to measure the fluid flow in the pipe.

An alternate location with sufficient straight piping was found, however this location did not capture all of the flow into the area. As a result some of the flow that supplies the space was not included in the measurement. The

flow to the Penthouse AHU and the four units located to the left of the riser in figure 8 was not included in the measurement. This is discussed in more detail later in this report.

For the control volume, measurements were taken from August 21, 2002 through August 23, 2002 and additional data was obtained from facility staff prior to and following the site monitoring. For this time of year, there were no heating loads.



The data collected is presented in Appendix A, and the analysis of this data is described in the following section.

Figure 10 Data Collection

ANALYSIS

Assumptions

The following assumptions were made to simplify the analysis and to keep irregularities in the data from corrupting the results of the study.

1. Electrical power measured at each study point was consumed on the fourth floor or by equipment supporting the fourth floor.
2. The power meter installed at Load Center 6 is accurate.

3. The power measured on PDP 4-1 is split evenly between the third and fourth floors.
4. The power measured on Load Center 6 is distributed evenly throughout the fourth floor.

Calculations

The data obtained through the three Dranetz-BMI power analyzers (recording power demand and consumption) was analyzed and averaged for a consistent 24-hour period of study. The same analysis was performed on the meter readings from the GE Load Center meter, the data from that analysis is provided in the tables below.

Power Demand			
Measurement Point	Demand (kW)		
	Min	Median	Max
PDP 4-1	75.45	76.42	87.02
CDP 4-4	190.33	191.69	194.47
PNL C-7	34.02	34.65	36.02
Load Center 6	N/A		

Table 4 Data Center Power Demand

Power Consumption					
Measurement Point	Start Time	Finish Time	Total Measurement Time	kWH	Avg 24 Hour kWH
PDP 4-1	8/21/02 16:30	8/23/02 8:30	40:00:00	3138.4	1883.04
CDP 4-4	8/21/02 16:30	8/23/02 8:15	39:45:00	7680.5	4637.28
PNL C-7	8/21/02 15:30	8/23/02 8:30	41:00:00	1433.2	838.95
Load Center 6	8/21/02 15:45	8/23/02 9:00	41:15:00	10800	6283.64

Table 5 Power Consumption

Power Distribution Panel 4-1 (PDP 4-1) feeds the computer room air conditioning (CRAC) units on the third and fourth floors. We assume that the power measured is evenly split between these two floors. This assumption seems reasonable due to the similarity in floor area. Based on this assumption, the measured power for an average twenty-four hour period by the 4th floor data center CRAC units is 941.5 kWH. This is the power necessary to power the fans and controls on the units. Energy to supply chilled water is not included here since its supply is from the central chilled water system.

Computer Distribution Panel 4-4 (CDP 4-4) supplies UPS power to the Power Distribution Units (PDU's) located in the fourth floor data center area. All of this power

supplies computer equipment located in the data center area. For an average twenty-four hour period the measured power consumption was 4637.3 kWh.

Load Center 6 (LC-6) feeds UPS power throughout the fourth floor and to Panel C-7 (PNL C-7) on the second floor. By measuring the energy used by LC-6 and deducting the energy used by PNL C-7 we can determine the UPS power used on the fourth floor. For an average twenty-four hour period this power consumption is 5344.7 kWh. Since this is power serving the whole fourth floor we assume that the power utilization is uniform across the floor and the data center under study comprises 15.7% of the fourth floor, therefore LC-6 supplied 837.6 kWh of energy to the control volume.

To get the total power consumed by the fourth floor the sum of the data for the adjusted Load Center 6, PDP 4-4, and CDP 4-1 was calculated. An equivalent amount of cooling must be supplied to the control volume to remove the heat. Table 6 summarizes these calculations:

4 th Floor Data Center Energy Consumption		
Total for 24 hour period	5614.8	kWH
Energy Consumed	19157712	BTU
Energy Flow Rate	798238	BTU/Hr
Energy Flow Rate	66.52	Tons of Refrigeration

Table 6 Control Volume total energy

In order to confirm the amount of chilled water being used to cool the fourth floor, we attempted to measure the flow and temperature differential of the chiller water feeding the fourth floor. When we arrived on site to install the fluid thermal energy meter, the intended site for installation was determined to be unsuitable for the connection of the meter. We chose an alternate location, however this secondary location did not capture all the flow to the fourth floor. The data collected by the meter is provided in Appendix A. A summary of this data is provided in table 7.

Total Flow	209,260 gal.
Avg Flow Rate	169.52 gpm
Avg dT	0.077°F
Period of Measure	41:15:00
Total Energy Flow	118,560 BTU
Avg 24 Hr Period Flow	170.68 gpm
Avg 24 Hr Energy Flow	70,694 BTU
Power Density	29.01 w/sf

Table 7 Control Volume Chilled Water Measurements

In order to determine if there is a significant relationship between weather patterns and data center power usage we gathered cooling degree days and heating degree day data for the time period under study and charted it against the power consumption for the site for the same time period. The data was taken from US Department of Commerce NOAA Historical Climatology data and the results are displayed below.

	Cooling Degree Days (65 deg F)	Heating Degree Days (65 deg F)	MWH
1997	569	5941	57463.2
1998	754	4948	42155.6
1999	875	5404	32174.7
2000	502	5910	32005.0
2001	713	5349	29861.2
2002	133	3224	31924.3

Table 8 Climatology data




Figure 11 Weather and Power Correlation

As the graph illustrates, there is no distinct correlation between the weather and the power usage by the site. This reinforces the 24 by 7 nature of data center operations and that the computing loads dominate regardless of the outdoor conditions.

Conclusions

CONTROL VOLUME B

The thermal energy measurements indicate that the CRAC systems being measured were not providing cooling. This may be the result of the areas on the fourth floor being idle or it could be due to the control of the two HVAC systems. Additionally our measurements did not capture the chilled water flow to the two air handlers in the penthouse above the fourth floor, these units apparently are handling the majority of the cooling for the fourth floor. It appears that the computer room air conditioning units are simply recirculating the air without providing cooling.

ENERGY EFFICIENCY OPPORTUNITIES:


GENERAL GUIDELINES FOR FUTURE CONSTRUCTION

Efficient Chilled Water Systems

Water cooled chillers offer enormous energy savings over air cooled chillers. Since the chiller is being cooled by lower temperature media, it can reject heat more easily, and does not have to work as hard. Though the addition of a cooling tower adds maintenance costs associated with the water treatment, we have found that the energy savings outweigh the maintenance costs. Within the options of water cooled chillers, variable

speed centrifugal are the most energy efficient, because they can operate very efficiently at low loads. The graph below compares the energy performance of various chiller types.

Comparison of Typical Chiller Efficiencies over Load Range



Chiller 1	250-Ton, Screw, Standard Efficiency, Air Cooled
Chiller 2	216 Ton, Screw, Water Cooled
Chiller 3	227-Ton, Centrifugal, Constant Speed, Water Cooled
Chiller 4	227-Ton, Centrifugal, Variable Speed, Water Cooled

Figure 12 Chiller Comparisons

Though there are efficient air cooled chillers, the larger size of water cooled chillers has resulted in more care given to efficiency and life cycle costs compared to air cooled chillers.

The selection of the auxiliary equipment, including the cooling tower, pumps, and pumping strategy should also be considered carefully. For example, variable speed fans on cooling towers allow for optimized cooling tower control. Premium efficiency motors and high efficiency pumps are recommended, and variable speed pumping is a ripe opportunity for pump savings. Variable pumping strategies can be achieved in a primary/secondary scheme, where the primary pumps operate at constant speed and directly feed water to the chiller, and the secondary pumps are variable speed and serve the air handling units. A more energy efficient scheme is primary-only variable speed

pumping strategy. Pumping savings are based on the cube law: pump power is reduced by the cube of the reduction in pump speed, which is directly proportional to the amount of fluid pumped.

A primary only variable pumping strategy must include a bypass valve that ensures minimum flow to the chiller, and the use of two-way valves at the air handling units in order to achieve lower pumping speeds. The control speed of the bypass valve should also meet the chiller manufacturers recommendations of allowable turndown, such that optimum chiller efficiency is achieved. This basically means that the flow through the chiller should be varied slow enough such that the chiller is able to reach a quasi-steady state condition and able to perform to its maximum efficiency. The diagram below describes the primary-only variable speed pumping strategy.

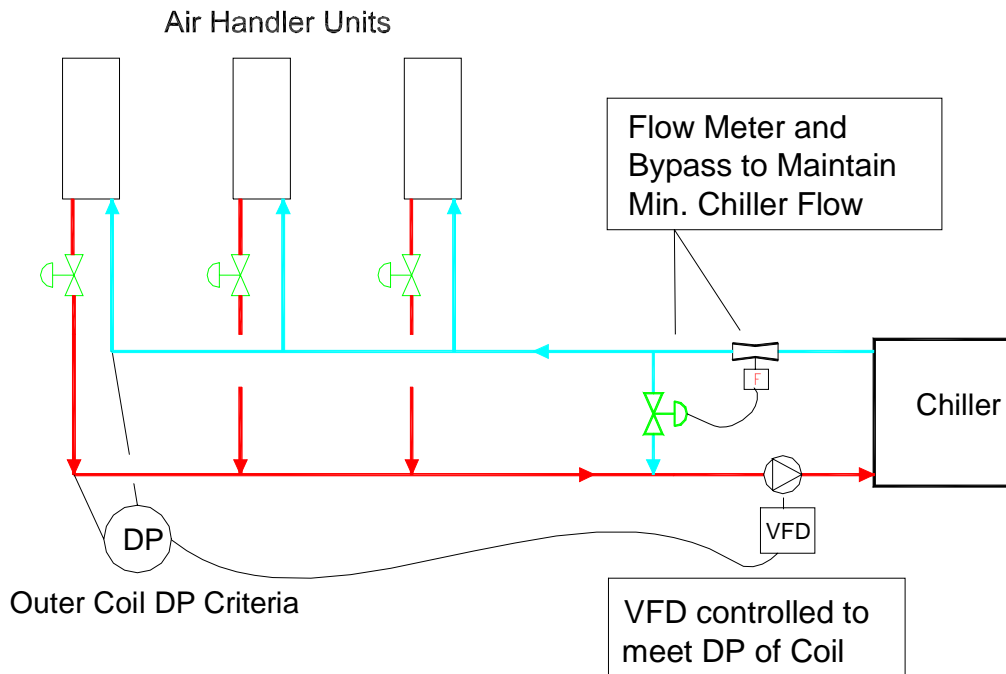


Figure 13 Air Handler Units

Air Management

A traditional method of cooling data centers employs an underfloor system fed by CRAC units. There are a number of potential problems with such systems: an underfloor system works on the basis of thermal stratification. This means that as the cool air is fed from the underfloor, it absorbs energy from the space, warming up as a result, and rises. In order to take advantage of thermal stratification, the return air must be collected at the ceiling level. CRAC units often have low return air grills, and are therefore simply recirculating cool or moderately warmed air. Furthermore, they are often located along the perimeter of the building, and not dispersed throughout the floor area where they can more effectively treat warm air. One alternative is to install transfer grills from the ceiling to the return grill.

Another common problem with underfloor supply is that the underfloor becomes congested with cabling, increasing the resistance to air flow. This results in an increase in fan energy use. A generous underfloor depth is essential for effective air distribution. Congested underfloor areas were observed in the control volume space as the following figure illustrates:



Figure 14 This under floor area in the NY Data Center severely restricts air flow

An alternative to underfloor air distribution is high velocity overhead supply, combined with ceiling height return. A central air handling system can be a very efficient air distribution unit. Design considerations include using VFDs on the fans, low pressure drop filters, and coils. An additional advantage of a central air handling system is that it can be specified with an economizer function.

Another common problem identified with CRAC units is that they are often fighting each other in order to maintain a constant humidity setpoint. Not only is a constant humidity setpoint unnecessary for preventing static electricity (the lower limit is more important), but it uses extra energy. A central air handling unit has better ability to control overall humidity than distributed CRAC units.

Air Management – Rack Configuration

Server rack configuration also dictates air management strategies in data centers. It is more logical for the aisles to be arranged such that servers' backs are facing each other, and servers' fronts are facing each other. This way, cool air is drawn in through the front, and hot air blown out the back (assuming a front to back server). The Uptime Institute has published documents describing this method for air management.

Commissioning of New Systems and Optimized Control Strategies

Many times the predicted energy savings of new and retrofit projects are not fully realized. Often, this is due to poor and/or incomplete implementation of the energy efficiency recommendations. Commissioning is the process of ensuring that the building systems perform as they were intended to by the design. Effective commissioning actually begins at the design stage, such that the design strategy is critically reviewed. Either the design engineer can serve as the commissioning agent, or a third party

commissioning agent can be hired. Commissioning differentiates from standard start-up testing in that it ensures systems function well relative to each other. In other words, it employs a systems approach.

Many of the problems identified in building systems are often associated with controls. A good controls scheme begins early in the design. In our experience, an effective controls design includes 1) a detailed points list, with accuracy levels, and sensor types, and 2) a detailed sequence of operations. Both of these components are essential for successfully implementing the recommended high efficiency chilled water system described above.

It is also possible that computer room air conditioners can be simultaneously cooling and humidifying – or heating and cooling at the same time. As noted below, however, it appears that cooling is not being provided by the CRAC units for the data center area examined.

Though use of commissioning is not uniformly adopted, various organizations have developed standards and guidelines. Such guidelines are available through organizations like Portland Energy Conservation Inc., at www.peci.org, or ASHRAE, Guideline 1-1996.

SPECIFIC OBSERVATIONS

Computer Room Air Conditioning

Based upon the measured data, it appears that very little, if any, cooling is provided by the computer room air conditioners. The inlet and exit temperatures of the chilled water to the CRAC units were approximately the same. It appears that the central building HVAC units are handling the entire cooling load for the computers and the CRAC units fan heat. If this is the case, there may also be greater outside air ventilation occurring than is necessary. Use of this system for total cooling may be the most efficient arrangement but deserves further study. The CRAC units are merely mixing the air in the data center areas. While this may be useful to avoid local hot spots, they apparently are not needed to provide overall cooling. It may be possible turn off some or all of the CRAC units and still achieve adequate cooling. This would save fan energy (and its added heat) for the units shut down. Further, with some modification to the distribution of the central system, it may be possible to eliminate the need for the CRAC units. Shutting off the units could save up to 330,400 kWh/year or \$25,000-\$30,000 per year.

The site should confirm that the central HVAC system is carrying the cooling load (by further monitoring of the chilled water inlet and outlet temperatures). If this is confirmed, the chilled water to the CRAC units could be then be turned off saving pumping energy. Approximately 169 gpm is provided to the CRAC units on a continuous basis. Given that the chilled water is being pumped to the fourth floor through

a long, complicated piping system, there is considerable opportunity for savings if the flow can be reduced or eliminated.

Openings were observed around many pieces of computing equipment through the raised floor. Sealing floor openings can improve efficiency by directing air through floor tiles to where it is needed. An air management scheme as described above should be followed by strategically placing floor tiles with openings in front of racks of computers.

Data center temperature and humidity control may be an opportunity for improvement. The observed temperature was colder than needed for human comfort and most likely lower than the electronics require. Studies have shown that the electrical components in data centers can withstand significantly higher temperatures and a broader range of humidity control, however the host site's customers would likely provide input to any change in operating criteria. It appears that the current computing loads are being accommodated through "brut force" cooling provided by the central air handlers. Temperatures throughout the data center could be monitored to confirm that local heat intensive areas are acceptable (if any exist).

Chilled Water System

Verify that the most efficient chiller is usually operating. Consider overall pumping energy for various combinations of chillers. Investigate use of free cooling and efficient operation of cooling towers. Various resources are available to provide guidance for chilled water systems, such as Cooltools:

(<http://www.hvacexchange.com/cooltools/coolhome.htm>)

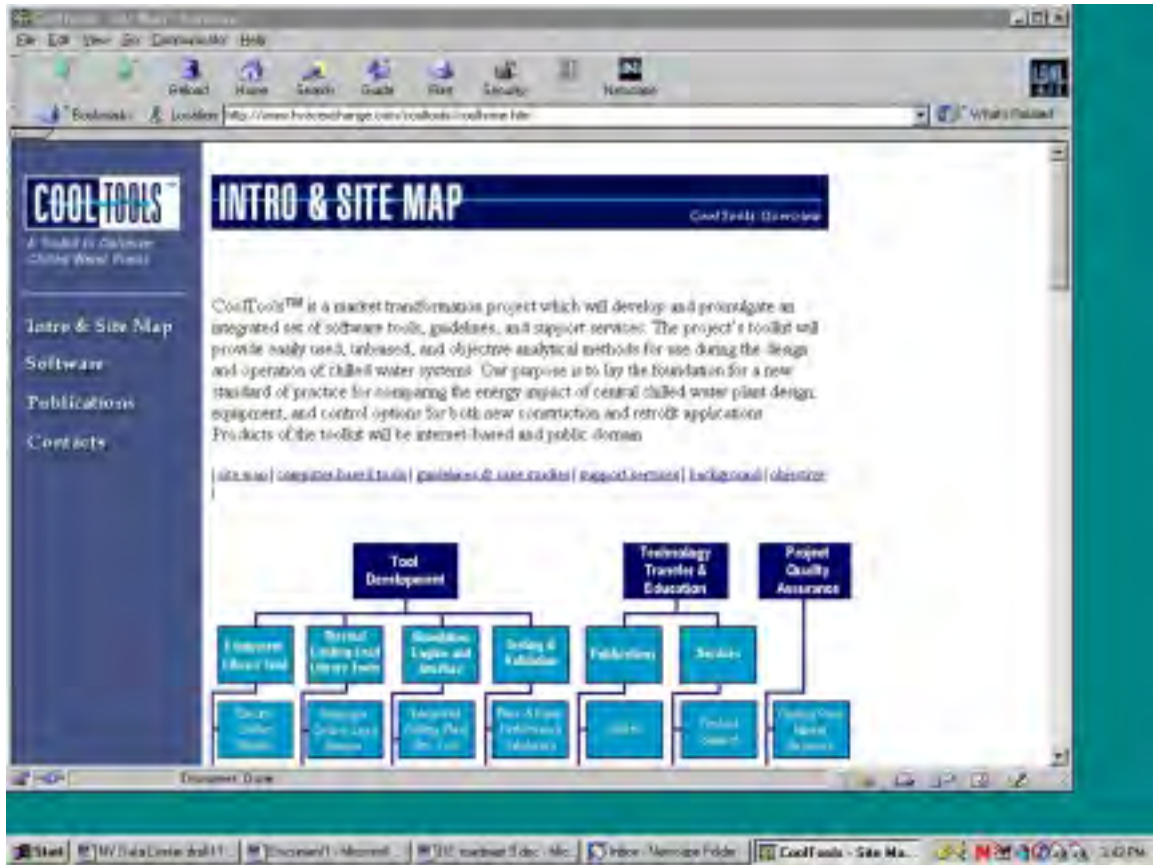


Figure 15 Cool Tools Web site

Computer room lighting

Consider use of standard lighting controls such as timers or occupancy sensors in data center areas. Many areas are unoccupied for long periods of time and comparable savings to office areas can be obtained. Consider reduced lighting levels and/or eliminating lighting in certain areas, especially in times of peak demand charges. Many telecom facilities hosting multiple customers are utilizing lighting controls to only illuminate a customer area when needed. This enhances energy savings and security. Savings for the direct cost of the lighting as well as the cost of removing the heat produced by the lighting will be realized.

Computing Equipment

Investigate ability to power down unused data processing equipment.

Appendices

Appendix A: Fluid Thermal Energy Meter Data (15 minute interval)

Appendix B: Power Monitoring Report

Appendix C: Data Center references

Appendix D: Monitoring Equipment

Appendix A: Fluid Thermal Energy Meter Data (15 minute interval)

Controlotron Data
15-minute intervals
Total: 41 hrs 15m

Date	Time	Heat Flow Rate MBTU/HR	Thermal Total MBTU	Flow Rate KGAL/MIN	Total Flow KGAL	Delta T dt(uS)
08.21.2002	15:45	0.21	0.09	0.17808	4.1	0.06821
08.21.2002	16:00	0.02	0.11	0.15772	6.45	0.07751
08.21.2002	16:15	-0.03	0.1	0.15904	8.86	0.07979
08.21.2002	16:30	-0.04	0.09	0.16341	11.31	0.0746
08.21.2002	16:45	-0.09	0.08	0.16533	13.76	0.07365
08.21.2002	17:00	-0.05	0.06	0.16393	16.19	0.0733
08.21.2002	17:15	-0.04	0.05	0.15955	18.62	0.0721
08.21.2002	17:30	-0.05	0.04	0.16692	21.08	0.07218
08.21.2002	17:45	-0.05	0.03	0.16048	23.49	0.08074
08.21.2002	18:00	-0.05	0.01	0.15777	25.88	0.07492
08.21.2002	19:30	-0.09	-0.08	0.16073	40.17	0.06786
08.21.2002	19:45	-0.08	-0.1	0.15737	42.55	0.07029
08.21.2002	20:00	-0.08	-0.12	0.15712	44.92	0.07952
08.21.2002	20:15	-0.08	-0.15	0.14388	47.2	0.06598
08.21.2002	20:30	-0.11	-0.17	0.15669	49.48	0.07263
08.21.2002	20:45	-0.09	-0.2	0.16235	51.89	0.07397
08.21.2002	21:00	-0.09	-0.22	0.16064	54.29	0.06896
08.21.2002	21:15	-0.1	-0.24	0.15805	56.67	0.0682
08.21.2002	21:30	-0.09	-0.26	0.15684	59.01	0.0677
08.21.2002	22:00	-0.09	-0.31	0.15903	63.68	0.07166
08.21.2002	22:15	-0.09	-0.33	0.1521	65.97	0.06885
08.21.2002	22:30	-0.1	-0.36	0.15255	68.25	0.06936
08.21.2002	22:45	-0.1	-0.38	0.14898	70.51	0.06286
08.21.2002	23:00	-0.11	-0.41	0.15238	72.75	0.0611
08.21.2002	23:15	-0.1	-0.44	0.14712	74.98	0.06834
08.21.2002	23:30	-0.1	-0.46	0.14592	77.21	0.06936
08.21.2002	23:45	-0.1	-0.49	0.14316	79.39	0.06633
08.22.2002	0:00	-0.11	-0.51	0.14438	81.55	0.07021
08.22.2002	0:15	-0.1	-0.54	0.14325	83.72	0.06607
08.22.2002	0:30	-0.1	-0.57	0.1411	85.85	0.06733
08.22.2002	0:45	-0.1	-0.59	0.14209	88.01	0.0572
08.22.2002	2:15	-0.11	-0.75	0.13863	100.66	0.066
08.22.2002	2:30	-0.06	-0.77	0.14309	102.75	0.07048
08.22.2002	2:45	-0.04	-0.78	0.1492	104.98	0.06838
08.22.2002	3:00	-0.04	-0.79	0.156	107.27	0.0712
08.22.2002	3:15	-0.03	-0.8	0.15522	109.6	0.07353

Date	Time	Heat Flow Rate MBTU/HR	Thermal Total MBTU	Flow Rate KGAL/MIN	Total Flow KGAL	Delta T dt(uS)
08.22.2002	3:30	-0.04	-0.8	0.15728	111.94	0.07286
08.22.2002	3:45	-0.04	-0.81	0.16208	114.33	0.07215
08.22.2002	4:00	-0.03	-0.82	0.16066	116.71	0.07533
08.22.2002	4:15	-0.03	-0.83	0.16193	119.12	0.07186
08.22.2002	4:30	-0.03	-0.83	0.16315	121.56	0.07406
08.22.2002	4:45	-0.02	-0.84	0.16504	124.03	0.07473
08.22.2002	5:00	-0.03	-0.85	0.16418	126.52	0.08268
08.22.2002	5:15	-0.03	-0.85	0.16583	129	0.07706
08.22.2002	5:30	-0.03	-0.86	0.16808	131.51	0.07232
08.22.2002	5:45	-0.03	-0.87	0.16693	134.01	0.08324
08.22.2002	6:00	-0.03	-0.87	0.1687	136.51	0.07396
08.22.2002	6:15	-0.03	-0.88	0.16568	139	0.07892
08.22.2002	6:30	-0.03	-0.89	0.16736	141.49	0.07264
08.22.2002	6:45	-0.03	-0.89	0.16788	144	0.07477
08.22.2002	7:00	-0.03	-0.9	0.16731	146.5	0.06875
08.22.2002	7:15	-0.03	-0.91	0.16613	149	0.07821
08.22.2002	7:30	-0.02	-0.91	0.16753	151.5	0.07436
08.22.2002	7:45	-0.01	-0.92	0.16752	154.05	0.07938
08.22.2002	8:00	-0.06	-0.93	0.16012	156.53	0.07161
08.22.2002	8:15	-0.09	-0.95	0.15151	158.83	0.06791
08.22.2002	8:30	-0.09	-0.97	0.15336	161.14	0.07168
08.22.2002	8:45	-0.09	-0.99	0.15439	163.46	0.0759
08.22.2002	9:00	-0.09	-1.02	0.15714	165.83	0.073
08.22.2002	9:15	-0.1	-1.04	0.1591	168.15	0.07045
08.22.2002	9:30	-0.09	-1.06	0.15812	170.54	0.07352
08.22.2002	9:45	-0.09	-1.08	0.16275	172.95	0.07239
08.22.2002	10:00	-0.09	-1.11	0.16564	175.39	0.07348
08.22.2002	10:15	-0.08	-1.13	0.1631	177.82	0.07523
08.22.2002	10:30	-0.1	-1.15	0.16385	180.24	0.06957
08.22.2002	10:45	-0.09	-1.18	0.16101	182.67	0.0762
08.22.2002	11:00	-0.09	-1.2	0.16635	185.16	0.07428
08.22.2002	11:15	-0.1	-1.22	0.16644	187.65	0.07455
08.22.2002	11:30	-0.09	-1.25	0.16928	190.16	0.07491
08.22.2002	11:45	-0.09	-1.27	0.16928	192.69	0.07544
08.22.2002	12:00	-0.08	-1.29	0.16761	195.22	0.0772
08.22.2002	12:15	-0.09	-1.31	0.17112	197.78	0.07605
08.22.2002	12:30	-0.06	-1.33	0.16929	200.32	0.07966
08.22.2002	12:45	-0.01	-1.34	0.17745	202.96	0.08341
08.22.2002	13:00	-0.01	-1.34	0.18218	205.69	0.08095
08.22.2002	13:15	0.01	-1.34	0.18514	208.44	0.08448
08.22.2002	13:30	0.02	-1.33	0.18496	211.19	0.08086
08.22.2002	13:45	0	-1.33	0.18296	213.95	0.08454
08.22.2002	14:00	0.01	-1.33	0.184	216.7	0.08084
08.22.2002	14:15	0.01	-1.33	0.18703	219.49	0.08865
08.22.2002	14:30	0.01	-1.33	0.18592	222.29	0.08231

Date	Time	Heat Flow Rate MBTU/HR	Thermal Total MBTU	Flow Rate KGAL/MIN	Total Flow KGAL	Delta T dt(uS)
08.22.2002	14:45	0.02	-1.32	0.18566	225.08	0.0825
08.22.2002	15:00	0.01	-1.32	0.1845	227.84	0.08217
08.22.2002	15:15	0.02	-1.32	0.18907	230.65	0.08599
08.22.2002	15:30	0.02	-1.31	0.18614	233.45	0.08499
08.22.2002	15:45	0.01	-1.31	0.18785	236.25	0.08622
08.22.2002	16:00	0.04	-1.3	0.18929	239.07	0.08356
08.22.2002	16:15	0.04	-1.29	0.18777	241.91	0.08598
08.22.2002	16:30	0.03	-1.29	0.1868	244.75	0.08544
08.22.2002	16:45	0.03	-1.28	0.18437	247.52	0.08986
08.22.2002	17:00	0.05	-1.26	0.18762	250.33	0.08397
08.22.2002	17:15	0.04	-1.25	0.18514	253.09	0.0815
08.22.2002	17:30	0.05	-1.24	0.18523	255.87	0.08461
08.22.2002	17:45	0.05	-1.23	0.18467	258.64	0.08081
08.22.2002	18:00	0.05	-1.22	0.18395	261.38	0.08397
08.22.2002	18:15	0.05	-1.21	0.18015	264.11	0.08146
08.22.2002	19:45	0.05	-1.12	0.18149	280.5	0.08071
08.22.2002	20:00	0.05	-1.11	0.18042	283.21	0.08223
08.22.2002	20:15	0.04	-1.1	0.18207	285.91	0.07933
08.22.2002	20:30	0.04	-1.09	0.1811	288.63	0.0808
08.22.2002	20:45	0.04	-1.08	0.18201	291.36	0.0819
08.22.2002	21:00	0.03	-1.07	0.17948	294.08	0.08414
08.22.2002	21:15	0.03	-1.06	0.17893	296.76	0.08185
08.22.2002	21:30	0.03	-1.05	0.18042	299.47	0.08672
08.22.2002	21:35	0.04	-1.05	0.18202	300.38	0.08451
08.22.2002	21:45	0.04	-1.04	0.18001	302.16	0.08468
08.22.2002	22:00	0.04	-1.03	0.17657	304.83	0.08178
08.22.2002	22:15	0.04	-1.02	0.17571	307.48	0.0783
08.22.2002	22:30	0.05	-1.01	0.17751	310.12	0.07788
08.22.2002	22:45	0.05	-1	0.17537	312.75	0.07877
08.22.2002	23:00	0.05	-0.99	0.17504	315.38	0.07741
08.22.2002	23:15	0.05	-0.97	0.17667	318.02	0.07922
08.22.2002	23:30	0.05	-0.96	0.17476	320.65	0.08074
08.22.2002	23:45	0.05	-0.95	0.17617	323.29	0.08192
08.23.2002	0:00	0.05	-0.93	0.17722	325.95	0.07759
08.23.2002	0:15	0.05	-0.92	0.17504	328.58	0.07872
08.23.2002	0:30	0.05	-0.91	0.17456	331.2	0.08119
08.23.2002	0:45	0.06	-0.89	0.1748	333.82	0.0829
08.23.2002	1:00	0.05	-0.88	0.17616	336.47	0.07901
08.23.2002	1:15	0.05	-0.86	0.17362	339.09	0.07759
08.23.2002	1:30	0.06	-0.85	0.17431	341.72	0.08262
08.23.2002	1:45	0.07	-0.83	0.17561	344.34	0.07911
08.23.2002	2:00	0.06	-0.82	0.17367	346.96	0.0763
08.23.2002	2:15	0.06	-0.8	0.17355	349.58	0.08122
08.23.2002	2:30	0.07	-0.79	0.17353	352.2	0.08103
08.23.2002	2:45	0.07	-0.77	0.17606	354.83	0.07651

Date	Time	Heat Flow Rate MBTU/HR	Thermal Total MBTU	Flow Rate KGAL/MIN	Total Flow KGAL	Delta T dt(uS)
08.23.2002	3:00	0.06	-0.76	0.17498	357.44	0.08108
08.23.2002	3:15	0.06	-0.74	0.17595	360.08	0.08255
08.23.2002	3:30	0.06	-0.72	0.17643	362.73	0.0826
08.23.2002	3:45	0.06	-0.71	0.17616	365.37	0.07669
08.23.2002	4:00	0.07	-0.69	0.17728	368.03	0.07914
08.23.2002	4:15	0.07	-0.68	0.17754	370.67	0.08242
08.23.2002	4:30	0.06	-0.66	0.17716	373.33	0.0845
08.23.2002	4:45	0.06	-0.64	0.17725	375.99	0.07967
08.23.2002	5:00	0.07	-0.63	0.17738	378.64	0.0778
08.23.2002	5:15	0.06	-0.61	0.17785	381.3	0.07825
08.23.2002	5:30	0.06	-0.6	0.17706	383.96	0.07938
08.23.2002	5:45	0.06	-0.58	0.17801	386.62	0.07935
08.23.2002	6:00	0.06	-0.57	0.17817	389.28	0.08251
08.23.2002	6:15	0.06	-0.55	0.17685	391.96	0.0815
08.23.2002	6:30	0.06	-0.54	0.17713	394.63	0.07781
08.23.2002	6:45	0.05	-0.52	0.17697	397.3	0.07979
08.23.2002	7:00	0.05	-0.51	0.1787	399.99	0.07984
08.23.2002	7:15	0.05	-0.5	0.17855	402.66	0.07686
08.23.2002	7:30	0.06	-0.48	0.17739	405.32	0.08144
08.23.2002	7:45	0.05	-0.47	0.18019	408.01	0.07691
08.23.2002	8:00	0.05	-0.46	0.17742	410.68	0.08327

Appendix B: Power Monitoring Report

Recovery Facility
KW PER SQ FT REPORT

DATE: AUGUST 28, 2002

Prepared for:

Jim McEnteggart
Syska & Hennessey

11 West 42nd Street
New York, NY 10001

Prepared By:

Paul Yetman, PE
SDM METRO

220 Maples Ave
Ph: 516-536-2600
Fax: 516-764-0755

-Table of Contents-

Executive Summary	pg. 2
PDP 4-1 KWH	pg. 3
PDP 4-1 Demand	pg. 4
PDP 4-1 Demand vs. Energy - 24 hours	pg. 5
PDP 4-1 Installation Photograph	pg. 6
CDP 4-4 KWH	pg. 7
CDP 4-4 Demand	pg. 8
CDP 4-4 Demand vs. Energy – 24 hours	pg. 9
CDP 4-4 Installation Photograph	pg. 10
Panel CP-7, Load Center 6 KWH	pg. 11
Panel CP-7, Load Center 6 Demand	pg. 12
Panel CP-7, Load Center 6 Demand vs. Energy – 24 hours	pg. 13
Panel CP-7, Load Center 6 Installation Photograph	pg. 14
Load Center 6: CDP 4-1, CDP 4-2, CDP 4-3, CP-7 kwh	pg. 15
Installation Photograph, Load Center 6-Meter	pg. 16

Executive Summary


On Wednesday, August 21, 2002 three Dranetz-BMI power analyzers were installed at the host site. The monitors were installed to determine the actual kW consumption per square foot of an operating data center on a daily basis. The monitors were installed at the following locations that fed the 4th floor of the data center: PDP 4-1, CDP 4-4, and Panel CP-7. Panel CP-7 does not supply power to the 4th floor data center but instead is fed from Load Center 6. Load Center 6 contained the following feeds that supplied power to the 4th Floor data center: CDP 4-1, CDP 4-2, and CDP 4-3.

The existing GE kWh meter was utilized as the measurement source for Load Center 6. The panel CP-7 was metered so this measurement could be subtracted from the GE kWh meter resulting in the power consumption total of CDP 4-1, CDP 4-2, and CDP 4-3.

This report provides the details of the power conditions found during the day and one half of monitoring.


PDP 4-1 - kwh

The graph below represents the minimum, maximum and median kwh consumption on PDP 4-1 over the entire monitoring period (1630 on 8/21 – 0830 on 8/23). The maximum number represents the total kwh consumed.




PDP 4-1 kw Demand

The graph below represents the minimum, maximum and median kw demand found on PDP 4-1 for the entire monitoring period (1630 on 8/21 – 0830 on 8/23).




PDP 4-1 Demand (kw) vs. Energy (kwh) - 24 hours

The graph below represents the minimum, maximum and median demand vs. energy conditions found PDP 4-1 on a 24-hour basis (1700 on 8/21 – 1700 on 8/22).




Installation Photograph – PDP 4-1

The photograph below shows the connection point of the Dranetz – BMI 4300 and associated current transformers on PDP 4-1.




CDP 4-4 kwh

The graph below represents the minimum, maximum and median kwh consumption on PDP 4-1 over the entire monitoring period (1630 on 8/21 – 0815 on 8/23). The maximum number represents the total kwh consumed.




CDP 4-4 kw Demand

The graph below represents the minimum, maximum and median kw demand found on PDP 4-1 for the entire monitoring period (1630 on 8/21 – 0815 on 8/23).




CDP4-4 Demand (kw) vs. Energy (kwh) - 24 hours

The graph below represents the minimum, maximum and median demand vs. energy conditions found PDP 4-1 on a 24-hour basis (1700 on 8/21 – 1700 on 8/22).




Installation Photograph – CDP 4-4

The photograph below shows the connection point of the Dranetz – BMI 4300 and associated current transformers on CDP 4-4.




Load Center 6 Panel CP-7 kwh

The graph below represents the minimum, maximum and median kwh consumption on CP-7 over the entire monitoring period (1530 on 8/21 – 0830 on 8/23). The maximum number represents the total kwh consumed.




Load Center 6 – Panel CP-7

The graph below represents the minimum, maximum and median kw demand found on Load Center 6, Panel CP-7 for the entire monitoring period (1530 on 8/21 – 0830 on 8/23)




Load Center 6, Panel CP-7 Demand (kw) vs. Energy (kwh)

The graph below represents the minimum, maximum and median demand vs. energy conditions found CP-7 on a 24-hour basis (1700 on 8/21 – 1700 on 8/22).



Installation Photograph – CP-7

The photograph below shows the connection point of the Dranetz – BMI 4300 and associated current transformers on CP-7.




Load Center 6 – CDP 4-1, CDP 4-2, CDP 4-3, CP-7 kwh

The GE meter which was recording the power consumption on Load Center 6 was reading 29379 @ 15:45 on 8/21/02. The reading was 29397 @ 9am on 8/23/02. The total change in the meter consumption for the entire period was 18 units. The meter had a multiplication factor of 600 resulting in a total consumption of 10,800 kwh for the entire load center over 41 hours and 15 minutes. The average consumption for the load center was 4.36 kwh per minute or 6,284 kwh for a 24-hour period.

Subtracting the total consumption of CP-7 of 860 kwh for a 24-hour period yields a 24 hour consumption of 5,424 kwh for CDP 4-1, CDP 4-2 and CDP 4-3.

Installation Photograph - Load Center 6 Meter

The photograph below shows the ending value for the metering on Load Center 6.



Appendix C: Data Center References

ACEEE, and CECS. 2001. Funding prospectus for "Analysis of Data Centers and their implications for energy demand". Washington, DC, American Council for an Energy Efficient Economy (ACEEE); Center for Energy and Climate Solutions (CECS). July 2001.

The paper includes an overview of data centers; discusses energy use, energy choices, and energy efficiency in data centers; potential impacts of data centers; present and future regulatory issues; and business opportunities in energy services.

Aebischer, B., R. Frischknecht, C. Genoud, A. Huser, and F. Varone. 2002a. Energy- and Eco-Efficiency of Data Centres. A study commissioned by Département de l'intérieur, de l'agriculture et de l'environnement (DIAE) and Service cantonal de l'énergie (ScanE) of the Canton of Geneva, Geneva, November 15.

The study investigates strategies and technical approaches to fostering more energy-efficient and environmentally sound planning, building and operating of data centres. It also formulate recommendations on how to integrate the findings in the legal and regulatory framework in order to handle construction permits for large energy consumers and promote energy efficiency in the economic sectors. Seventeen recommendations grouped in four topics are derived from study conclusions: Transfer of the accord into an institutionalised legal and regulatory framework; Energy-efficiency policies for all large energy consumers; Preconditions, and prerequisites; Operational design of voluntary energy policies.

Aebischer, B., R. Frischknecht, C. Genoud, and F. Varone. 2002b. Energy Efficiency Indicator for High Electric-Load Buildings. The Case of Data Centres. Proceedings of the IEECB 2002. 2nd International Conference on Improving Electricity Efficiency in Commercial Buildings. Nice, France.

Energy per unit of floor area is not an adequate indicator for energy efficiency in high electric-load buildings. For data centres we propose to use a two-stage coefficient of energy efficiency $CEE = C1 * c2$, where $C1$ is a measure of the efficiency of the central infrastructure and $c2$ a measure of the energy efficiency of the equipment.

Anonymous. 2001. Model Data Center Energy Design Meeting. Austin Energy, Austin, TX, Feb 12-13.

http://www.austinenergy.com/business/energy_design_meeting.htm

Anonymous. 2002a. 7 x 24 Update: Design & Construction - Issues and trends in mission critical infrastructure design, planning and maintenance.

<http://www.facilitiesnet.com/BOM/Jan02/jan02construction.shtml>. July 23, 2002.

<http://www.7x24exchange.org/>.

- Anonymous. 2002b. Continuous Availability Review (CAR). The Uptime Institute: Computersite Engineering, Inc. <http://www.upsite.com/csepages/csecar.html>. July 22, 2002.
- Anonymous. 2002c. End-to-End Reliability Begins with the User's Definition of Success. The Uptime Institute. <http://www.upsite.com/TUIpages/editorials/endtoend.html>. July 22, 2002.
- Anonymous. 2002d. Mechanical Systems Diagnostic Review (MSDR). The Uptime Institute: Computersite Engineering, Inc. <http://www.upsite.com/csepages/csemsdr.html>. July 22, 2002.
- Anonymous. 2002e. Site Infrastructure Operations Review (SIOR). The Uptime Institute: Computersite Engineering, Inc. <http://www.upsite.com/csepages/cseior.html>. July 22, 2002.
- Baer, D. B. Emerging Cooling Requirements & Systems in Telecommunications Spaces, Liebert Corporation.
During the last several years, power density trends, and consequently thermal density trends in telecommunications spaces have become topics of increasing interest. This paper identifies several of the underlying drivers of these trends, project possible outcomes, and assess the impact on cooling system design for these spaces.
- Beck, F. 2001. Energy Smart Data Centers: Applying Energy Efficient Design And Technology To The Digital Information Sector. *Renewable Energy Policy Project (REPP): Washington, DC*. (November 2001 REPP).
Both utilities and data center owners face challenges in meeting electricity demand loads with required levels of reliability. However, the bursting of the high-tech stock bubble in 2000 and the 2001 U.S. economic downturn has slowed expansion of data centers. This provides time and an opportunity to examine data center construction and operational practices with an eye toward reducing their energy demands through use of energy efficient technologies and energy smart design practices. As the economy recovers and the next data center rush approaches, best practices can reduce energy use while maintaining or even increasing data center reliability. Energy demands of data centers that support the digital information- and communications-based economy need not be as high as some predict. In fact, data center power demands could be reduced by 20 percent with minimal efficiency efforts, and by 50 percent with more aggressive efficiency measures.
- Blount, H. E., H. Naah, and E. S. Johnson. 2001. Data Center and Carrier Hotel Real Estate: Refuting the Overcapacity Myth. Lehman Brothers: TELECOMMUNICATIONS, New York, June 7, 2001. <http://www.lehman.com>
An exclusive study examining supply and demand trends for data center and carrier hotel real estate in North America. Lehman Brothers and Cushman &

Wakefield have completed the first in a regular series of proprietary studies on telecommunications real estate (TRE), including carrier hotels and data centers.

Bors, D. 2000. Data centers pose serious threat to energy supply. *Puget Sound Business Journal (Seattle)* - October 9, 2000.

<http://seattle.bizjournals.com/seattle/stories/2000/10/09/focus5.html>

To cope with increasing energy demand from data centers, the author discussed feasibilities of two possible approaches: 1) energy industry approach by looking at alternative energy supply; 2) construction industry approach by looking at data center energy efficiency. To get there, it is worth investigating four distinct components. (I) Co-generation of power. Presently, standby diesel generators are required to maintain the desired level of reliability at most data center sites, but their exhaust makes most of these generators unacceptable for long-term power generation. (II) Fuel cells offer the promise of very clean emissions and the reasonable possibility for use as standby power. (III) Increased efficiency in data center power distribution systems. There are two separate items that are major contributors to data center power distribution system inefficiencies. The first, power distribution units (PDUs), are available with optional internal transformers that use less energy than the present cadre of K-rated transformers. The second, uninterruptible power systems (UPSs), come in a range of efficiency ratings. If the use of high-efficiency PDUs and UPSs are combined, they offer the potential of a 6 percent saving. (IV) Increased efficiency in mechanical cooling systems. In order to ensure data center reliability, mechanical equipment is often selected as a large number of small, self-contained units, which offers opportunities to improve efficiencies. (V) Reductions in energy use by computer, network and storage equipment. Computer manufacturers can do their part by creating computers with greater computational power per watt. They have been doing this for years as a side effect of hardware improvements, and they can do even better if they make it a goal.

Brown, E., R. N. Elliott, and A. Shipley. 2001. Overview of Data Centers and Their Implications for Energy Demand. Washington, DC, American Council for an Energy Efficient Economy, Center for Energy & climate Solutions (CECS). September 2001. <http://www.aceee.org/pdfs/datacenter.pdf.pdf>

The white paper discusses data center industry boom and energy efficiency opportunities and incentives in internet data centers. Emerging in the late 1990's, data centers are locations of concentrated Internet traffic requiring a high-degree of power reliability and a large amount of power relative to their square footage. Typically, power needs range from 10-40MW per building, and buildings are typically built in clusters around nodes in the Internet fiber-optic backbone. During the development boom in 1999 and 2000, projects averaged 6-9 months from site acquisition to operation, and planned operational life was 36 months to refit. Even high energy-prices were dwarfed by net daily profits of 1-2 million dollars per day for these buildings during the boom, creating little incentive for efficient use of energy.

- Callsen, T. P. 2000. The Art of Estimating Loads. *Data Center* (Issue 2000.04).
This article discusses the typical Data Center layout. It includes floor plan analysis, HVAC requirements, and the electrical characteristics of the computer hardware typically found in a Data Center.
- Calwell, C., and T. Reeder. 2002. Power Supplies: A Hidden Opportunity for Energy Savings (An NRDC Report). Natural Resources Defense Council, San Francisco, CA, May 22, 2002. <http://www.nrdc.org>
The article discusses the efficiency of power supplies which perform current conversion and are located inside of the electronic product (internal) or outside of the product (external). The study finds that most external models, often referred to as "wall-packs" or "bricks," use a very energy inefficient design called the linear power supply, with measured energy efficiencies ranging from 20 to 75%; that most internal power supply models use somewhat more efficient designs called switching or switch-mode power supplies; and that internal power supplies have energy efficiencies ranging from 50 to 90%, with wide variations in power use among similar products. Most homes have 5 to 10 devices that use external power supplies, such as cordless phones and answering machines. Internal power supplies are more prevalent in devices that have greater power requirements, typically more than 15 watts. Such devices include computers, televisions, office copiers, and stereo components. The paper points out that power supply efficiency levels of 80 to 90% are readily achievable in most internal and external power supplies at modest incremental cost through improved integrated circuits and better designs.
- Cratty, W., and W. Allen. 2001. Very High Availability (99.9999%) Combined Heat and Power for Mission Critical Applications. *Cinintel 2001*: 12.
<http://www.surepowersystem.com>
- Elliot, N. 2001. Overview of Data Centers and their implications for energy demand. Washington, DC, American Council for an Energy Efficient Economy. Jan 2001, revised June 10, 2001.
- Feng, W., M. Warren, and E. Weigle. 2002. The Bladed Beowulf: A Cost-Effective Alternative to Traditional Beowulfs. *Cluster2002 Program*. <http://www-unix.mcs.anl.gov/cluster2002/schedule.html>; public.lanl.gov/feng/Bladed-Beowulf.pdf
Authors present a novel twist to the Beowulf cluster - the Bladed Beowulf. In contrast to traditional Beowulfs which typically use Intel or AMD processors, the Bladed Beowulf uses Transmeta processors in order to keep thermal power dissipation low and reliability and density high while still achieving comparable performance to Intel- and AMD-based clusters. Given the ever increasing complexity of traditional super-computers and Beowulf clusters; the issues of size, reliability, power consumption, and ease of administration and use will be "the" issues of this decade for high-performance computing. Bigger and faster machines are simply not good enough anymore. To illustrate, Authors present the

results of performance benchmarks on the Bladed Beowulf and introduce two performance metrics that contribute to the total cost of ownership (TCO) of a computing system - performance/power and performance/space.

Frith, C. 2002. Internet Data Centers and the Infrastructure Require Environmental Design, Controls, and Monitoring. *Journal of the IEST* **45**(2002 Annual Edition): 45-52.

Internet Data Centers and the Infrastructure Require Environmental Design, Controls, and Monitoring. The author points out that specifications and standards need to be developed to achieve high performance for mission-critical internet applications.

Gilleskie, R. J. 2002. The Impact of Power Quality in the Telecommunications Industry. Palm Springs, CA, June 4. <http://www.energy2002.ee.doe.gov/Facilities.htm>

The workshop addresses the unique issues and special considerations necessary for improving the energy efficiency and reliability of high-tech data centers. This presentation addresses impacts of power quality including voltage sags, harmonics, and high-frequency grounding in telecommunication industry.

Grahame, T., and D. Kathan. 2001. Internet Fuels Shocking Load Requests. *Electrical World* **Vol. 215** (3): 25-27.

http://www.platts.com/engineering/ew_back_issues.shtml

This article discusses the implications of the increase for power demand by the Internet's traffic growth on utility planning, operation, and financing.

Greenberg, D. 2001. Addendum to ER-01-15: A Primer on Harmonics. E-SOURCE, Boulder, Colorado, September 2001.

The electrical distribution systems of most commercial and industrial facilities were not designed to operate with an abundance of harmonics-producing loads. In fact, it is only within recent years that such loads have become widespread enough for industry to take notice and to begin to develop strategies to address the problems that harmonics can create. By 1992, concern about the issue had grown sufficiently that the Institute for Electrical and Electronic Engineers (IEEE) developed and published its standard 519, "IEEE Recommended Practices and Requirements for Harmonic Control in Electrical Power Systems," which established an approach for setting limits on the harmonic voltage distortion on the utility power system and on the harmonic currents created individual power consumers. Since that time, the electronic loads that give rise to harmonic currents have grown dramatically and are projected to continue growing for the foreseeable future. This being the case, there is and will continue to be a market for technological solutions to the problems that harmonics can cause.

Gross, P. 2002. Needed: New Metrics. *Energy User News*.

http://www.energyusernews.com/eun/cda/articleinformation/features/bnp__features__item/0,2584,82741,00.html

- Gruener, J. 2000. Building High-Performance Data Centers. *Dell Magazines - Dell Power Solutions* (Issue 3 "Building Your Internet Data Center").
http://www.dell.com/us/en/esg/topics/power_ps3q00_1_power.htm;
http://www.dell.com/us/en/esg/topics/power_ps3q00-giganet.htm
 The introduction of Microsoft SQL Server 2000 is a milestone in the race to build the next generation of Internet data centers. These new data centers are made up of tiers of servers, now commonly referred to as server farms, which generally are divided into client services servers (Web servers), application/business logic servers, and data servers supporting multiple instances of databases such as SQL Server 2000.
- Hellmann, M. 2002. Consultants Face Difficult New Questions in Evolving Data Center Design. *Energy User News*.
http://www.energyusernews.com/CDA/ArticleInformation/features/BNP__Features__Item/0,2584,70610,00.html
 While few data center design projects are alike, there are always the twin challenges of "power and fiber." And sometimes, even local politics and human factors. The paper suggested that the consultant should be brought in as soon as a business case is established so criteria can be established and a concept can be developed, priced, and compared to the business case. A planning is necessary before moving on to site selection and refine the concept and again test the business case.
- Howe, B., A. Mansoor, and A. Maitra. 2001. Power Quality Guidelines for Energy Efficient Device Application - Guidebook for California Energy Commission (CEC). Final Report to B. Banerjee, California Energy Commission (CEC).
 Energy efficiency and conservation are crucial for a balanced energy policy for the Nation in general and the State of California. Widespread adaptation of energy efficient technologies such as energy efficient motors, adjustable speed drives, improved lighting technologies will be the key in achieving self sufficiency and a balanced energy policy that takes into account both supply side and demand side measures. In order to achieve the full benefit of energy efficient technologies, these must be applied intelligently, and with clear recognition of the impacts some of these technologies may have on power quality and reliability. Any impediment to the application of these energy efficient technologies by the customers is not desirable for the overall benefit to energy users in California. With that in mind EPRI and CEC has worked to develop this guidebook to promote customer adaptation of energy efficient technologies by focusing on three distinct objectives. 1) Minimize any undesirable power quality impacts of energy-saving technologies; 2) Understand the energy savings potential of power quality-related technologies. These include: Surge Protective Devices (SPDs) or Transient Voltage Surge Suppressors (TVSS), Harmonic Filters, Power Factor Correction Capacitors, Electronic Soft Starters for Motors; and 3) How to evaluate "black box" technologies

Intel. 2002. Planning and Building a Data Center - Meeting the e-Business Challenge. Intel Corp.

http://www.intel.com/network/idc/doc_library/white_papers/data_center/. Aug 01, 2002.

The paper discusses the keys to success of Internet Service Providers (ISPs) that include 1) Achieve the economies of scale necessary to support a low price business model; 2) Offer added value, typically in the form of specialized services such as applications hosting to justify a premium price. This document provides a high-level overview of the requirements for successfully establishing and operating an Internet data center in today's marketplace. It offers some of the key steps that need to be taken, including project definition, prerequisites and planning. In order to construct a data center that can meet the challenges of the new market, there are three basic areas of data center definition and development: 1) Facilities: including building, security, power, air-conditioning and room for growth; 2) Internet connectivity: performance, availability and scalability; 3) Value-added services and the resources to support their delivery: service levels, technical skills and business processes. The aim is to provide customers with the physical environment, server hardware, network connectivity and technical skills necessary to keep Internet business up and running 24 hours a day, seven days a week. The ability to scale is essential, allowing businesses to upgrade easily by adding bandwidth or server capacity on demand.

Koplin, E. 2000. Finding Holes In The Data Center Envelope. *Engineered Systems* (September 2000).

http://www.esmagazine.com/CDA/ArticleInformation/features/BNP__Features__Item/0,2503,8720,00.html

The paper addresses importance of environmental control in data center facilities. Maintaining data center availability requires absolutely reliable infrastructure. A significant amount of this is devoted solely to maintaining stable environmental parameters. And only constant, thorough regulation and testing of these parameters ensures the integrity of the data center "envelope."

Mandel, S. 2001. Rooms that consume - Internet hotels and other data centers inhale electricity. *Electric Perspectives* Vol. 26 (No.3).

http://www.eei.org/ep/editorial/Apr_01/0401ROOM.htm

The article estimated that the amount of this data center space in the United States nearly doubled in 2000, totaling between 19 million and 25 million square feet by year-end, according to investment analysts. They say they expect another 10 million to 20 million square feet of new space to be added in 2001. Developers are asking electric utilities to supply the buildings with 100-200 watts of electricity per square foot. Since these data centers are new to the economy, there is little historical data on which to base estimates of electricity use for a facility. In addition, the dot.com world makes it difficult for the developer to say confidently how much electricity one of these internet hotels will use. Source One estimates that tens of billions of dollars worth of electric infrastructure improvements will be needed for data centers over the next few years and that

they will consume billions of dollars more worth of electricity. The energy costs are as high or higher than the actual lease costs. Indeed, 50-60 percent of the cost of building a data center is for the power, including batteries, backup generators, and air-conditioning, as well as the cost for utility construction.

Mitchell-Jackson, J. 2001. Energy Needs in an Internet Economy: A Closer Look at Data Centers, July, 2001.

This study explains why most estimates of power used by data centers are significantly too high, and gives measured power use data for five such facilities. Total power use for the computer room area of these data centers is no more than 40 W/square foot, including all auxiliary power use and cooling energy. There are two draft journal articles from this work, one focusing on the detailed power use of the data center we've examined in most detail, and the other presenting the aggregate electricity use associated with hosting-type data centers in the U.S.

Mitchell-Jackson, J., J. G. Koomey, B. Nordman, and M. Blazek. 2001. Data Center Power Requirements: Measurements From Silicon Valley. *Energy—the International Journal* (Under review).

<http://enduse.lbl.gov/Projects/InfoTech.html>

Current estimates of data center power requirements are greatly overstated because they are based on criteria that incorporate oversized, redundant systems, and several safety factors. Furthermore, most estimates assume that data centers are filled to capacity. For the most part, these numbers are unsubstantiated. Although there are many estimates of the amount of electricity consumed by data centers, until this study, there were no publicly available measurements of power use. This paper examines some of the reasons why power requirements at data centers are overstated and adds actual measurements and the analysis of real-world data to the debate over how much energy these facilities use.

Patel, C. D., C. E. Bash, C. Belady, L. Stahl, and D. Sullivan. 2001. Computational Fluid Dynamics Modeling of High Compute Density Data Centers to Assure System Inlet Air Specifications. Reprinted from the proceedings of the Pacific Rim ASME International Electronic Packaging Technical Conference and Exhibition (IPACK 2001), © 2001, ASME.

Due to high heat loads, designing the air conditioning system in a data center using simple energy balance is no longer adequate. Data center design cannot rely on intuitive design of air distribution. It is necessary to model the airflow and temperature distribution in a data center. This paper presents a computational fluid dynamics model of a prototype data center to make the case for such modeling.

Patel, C. D., R. Sharma, C. E. Bash, and A. Beitelmal. 2002. Thermal Considerations in Cooling Large Scale High Compute Density Data Centers. 8th IThERM Conference. San Diego CA.

A high compute density data center of today is characterized as one consisting of thousands of racks each with multiple computing units. The computing units

include multiple microprocessors, each dissipating approximately 250 W of power. The heat dissipation from a rack containing such computing units exceeds 10 KW. Today's data center, with 1000 racks, over 30,000 square feet, requires 10 MW of power for the computing infrastructure. A 100,000 square foot data center of tomorrow will require 50 MW of power for the computing infrastructure. Energy required to dissipate this heat will be an additional 20 MW. A hundred thousand square foot planetary scale data center, with five thousand 10 KW racks, would cost ~\$44 million per year (@ \$100/MWh) just to power the servers & \$18 million per year to power the cooling infrastructure for the data center. Cooling design considerations by virtue of proper layout of racks can yield substantial savings in energy. This paper shows an overview of a data center cooling design and presents the results of a case study where layout change was made by virtue of numerical modeling to avail efficient use of air conditioning resources.

PG&E. 2001. Data Center Energy Characterization Study. Pacific Gas and Electric Company (subcontractor: Rumsey Engineers), San Francisco, Feb. 2001. Rumsey Engineers, Inc. and PG&E have teamed up to conduct an energy study as part of PG&E's Data Center Energy Characterization Study. This study will allow PG&E and designers to make better decisions about the design and construction of data centers in the near future. Three data centers in the PG&E service territory have been analyzed during December 2000 and January 2001, with the particular aim of determining the end-use of electricity. The electricity use at each facility was monitored for a week each. At the end of the report are a set of definitions, which explain the terms used and the components in making each calculation. The three data centers provide co-location service, which is an unmanaged service that provides rack space and network connectivity via a high capacity backbone. About half or more of the electricity goes to powering the data center floor, and 25 to 34 percent of the electricity goes to the heating, air conditioning and ventilation equipment. The HVAC equipment uses a significant amount of power and is where energy efficiency improvements can be made. All three facilities use computer room air conditioning (CRAC) units, which are stand-alone units that create their own refrigeration and circulate air. A central, water-cooled chilled water system with air handlers and economizers can provide similar services with roughly a 50% reduction in cooling energy consumption. Energy density of the three buildings had an average of 35 W/sf. The cooling equipment energy density for the data center floor alone averaged at 17 W/sf for the three facilities. The average designed energy density of the three data centers' server loads was 63 W/sf, while the measured energy density was 34 W/sf. An extrapolated value was also calculated to determine what the server load energy density would be when fully occupied. The average extrapolated energy density was 45 W/sf. Air movement efficiency varies from 23 to 64 percent between the three facilities. Cooling load density varies from 9 to 70 percent between the three facilities.

Planet-TECH. 2002. Technical and Market Assessment for Premium Power in Haverhill. Planet-TECH Associates for The Massachusetts Technology Collaborative,

www.mtpc.org, Westborough, MA 01581-3340, Revision: February 20, 2002.
http://www.mtpc.org/cluster/Haverhill_Report.pdf ; <http://www.planet-tech.com/content.htm?cid=2445>

This study is pursued under contract to the Massachusetts Technology Collaborative, in response to a request for a "Technical and Market Assessment". It seeks to determine if the provisioning of "premium power" suitable for data-intensive industries will improve the marketability of a Historic District mill building in Haverhill. It is concluded that such provisioning does improve the marketability, however, not to a degree that is viable at this time. Other avenues for energy innovation are considered and recommendations for next steps are made.

RMI, and DR International. 2002. Energy Efficient Data Centers - A Rocky Mountain Institute Design Charrette. Organized, Hosted and Facilitated by Rocky Mountain Institute, with D&R International, Ltd. and Friends. Hayes Mansion Conference Center, San Jose, California.

Rapid growth of "mission critical" server-farm and fiber-optic-node data centers has presented energy service providers with urgent issues. Resulting costs have broad financial and societal implications. While recent economic trends have severely curtailed projected growth, the underlying business remains vital. The current slowdown allows us all some breathing room—an excellent opportunity to step back and carefully evaluate designs in preparation for surviving the slowdown and for the resumption of explosive growth. Future data center development will not occur in the first-to-market, damn-the-cost environment of 1999-2000. Rather, the business will be more cost-competitive, and designs that can deliver major savings in both capital cost (correct sizing) and operating cost (high efficiency)—for both new build and retrofit—will provide their owners and operators with an essential competitive advantage.

Robertson, C., and J. Romm. 2002. Data Centers, Power, and Pollution Prevention - Design for Business and Environmental Advantage. The Center for Energy and Climate Solutions; A Division of The Global Environment and Technology Foundation, June 2002. <http://www.cool-companies.org>; <http://www.getf.org>
Computers and other electronic equipment will crash at the slightest disruption or fluctuation in their supply of electricity. The power system was not designed for these sensitive electronic loads and is inherently unable to meet the technical requirements of the information economy. For data centers, which play a central role in the information economy, crashing computers cause potentially catastrophic financial losses. The same voltage sag that causes the lights to dim briefly can cause a data center to go off-line, losing large sums of money, for many hours. Data center owners and their power providers must therefore solve several related technical and economic electric power problems. These are: 1) How to assure high-availability (24x7) power supply with a very low probability of failure; 2) How to assure practically perfect power quality; and 3) How to manage risk while minimizing capital and operating expenses

Roth, K. W., Fred Goldstein, and J. Kleinman. 2002. Energy consumption by office and telecommunications equipment in commercial buildings, Volume I: Energy Consumption Baseline. Arthur D. Little (ADL), Inc., 72895-00, Cambridge, MA, January 2002.

ADL carried out a "bottom-up" study to quantify the annual electricity consumption (AEC) of more than thirty (30) types of non-residential office and telecommunications equipment. A preliminary AEC estimate for all equipment types identified eight key equipment categories that received significantly more detailed studied and accounted for almost 90% of the total preliminary AEC. The Key Equipment Categories include: *Computer Monitors and Displays, Personal Computers, Server Computers, Copy Machines, Computer Network Equipment, Telephone Network Equipment, Printers, Uninterruptible Power Supplies (UPSs)*. The literature review did not uncover any prior comprehensive studies of telephone network electricity consumption or uninterruptible power supply (UPS) electricity consumption. The AEC analyses found that the office and telecommunications equipment consumed 97-TWh of electricity in 2000. The report concludes that commercial sector office equipment electricity use in the U.S. is about 3% of all electric power use. The ADL work also creates scenarios of future electricity use for office equipment, including the energy used by telecommunications equipment.

Sullivan, R. F. 2002. Alternating Cold and Hot Aisles Provides More Reliable Cooling for Server Farms. The Uptime Institute.

<http://www.uptimeinstitute.org/tuiaisles.html>

The creation of "server farms" comprising hundreds of individual file servers has become quite commonplace in the new e-commerce economy, while other businesses spawn farms by moving equipment previously in closets or under desktops into a centralized data center environment. However, many of these farms are hastily planned and implemented as the needed equipment must be quickly installed on a rush schedule. The typical result is a somewhat haphazard layout on the raised floor that can have disastrous consequences due to environmental temperature disparities. Unfortunately, this lack of floor-layout planning is not apparent until after serious reliability problems have already occurred.

The Uptime Institute. 2000. Heat-Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment. The Uptime Institute, Version 1.0., <http://www.upsite.com/>. <http://www.uptimeinstitute.org/heatdensity.html>

This white paper provides data and best available insights regarding historical and projected trends in power consumption and the resulting heat dissipation in computer and data processing systems (servers and workstations), storage systems (DASD and tape), and central office-type telecommunications equipment. The topics address the special needs of Information Technology professionals, technology space and data center owners, facilities planners, architects, and engineers.

Thompson, C. S. 2002. Integrated Data Center Design in the New Millennium. *Energy User News*.

http://www.energyusernews.com/CDA/ArticleInformation/features/BNP_Features_Item/0,2584,70578,00.html

Data center design requires planning ahead and estimating future electrical needs. Designers must accurately predict space and energy requirements, plus cooling needs for new generations of equipment. Importance of data center reliability is discussed.

Wood, L. 2002. Cutting Edge Server Farms - The blade server debate. [newarchitectmag.com](http://www.newarchitectmag.com).

<http://www.newarchitectmag.com/documents/s=2412/na0702f/index.html>. July 23, 2002. A blade is the industry term for a server that fits on a single circuit board, including CPU, memory, and perhaps a local hard disk. Multiple blades are plugged into a chassis, where each blade shares a common power supply, cooling system, and communications back plane. Multiple chassis can then be stacked into racks. By comparison, the conventional approach for rack-mounted servers involves only one server per chassis. A chassis cannot be smaller than one vertical rack unit (1U, or about 1.75 inches high). This limits you to 42 to 48 servers in a standard seven-foot rack. A typical blade chassis is much higher than 1U, but several can still be stacked in a rack, allowing upwards of 300 servers per rack, depending on the vendor and configuration. This compact design offers compelling advantages to anyone operating a high-density server farm where space is at a premium. Indeed, blades are the "next big thing" in servers, and it's probable that any given administrator will have to decide whether to adopt them in the near future.

Appendix XII. “NY Data Center Energy Benchmarking and Case Study, Facility 11”; NYSERDA Case Study, Syska & Hennessy; LBNL – William Tschudi, 2003

NY DATA CENTER No. 2 ENERGY BENCHMARKING AND CASE STUDY



JULY, 2003

**PRINCIPAL INVESTIGATOR:
SUBCONTRACTOR:**

**WILLIAM TSCHUDI - LBNL
PAUL LIESMAN - SYSKA HENNESSY GROUP**

I. Executive Summary

The New York State Energy Research and Development Authority (NYSERDA) sponsored this project to study energy use in New York data centers. This facility is the second New York case study along with over a dozen case studies performed in California. Each case study focuses on energy efficiency improvement possibilities, establishes benchmarks for key metrics, and determines energy end use. The results are reported anonymously for each of the case studies and this report is termed Facility 14. Additional case studies and benchmark results are provided on LBNL's website (<http://datacenters.lbl.gov>)

The information obtained in these studies provides insight into the electrical power use within the data center and the overall electrical demand that data center facilities create for utilities. Energy benchmark data for a larger number of data centers will further help to identify current best practices, and point to efficiency and reliability improvement areas.

Facility 14 utilizes typical computer room air conditioning (CRAC) units for cooling the computing equipment (figure 1). The CRAC units draw heated return air in from the top, and provide conditioned air under a raised floor. The air is then distributed directly to the computing equipment cabinets through openings in the raised floor. Fan energy was obtained for all CRAC units as one of the energy end uses for the data center.



Figure 1 CRAC Unit



Figure 2 Centrifugal Chiller

Chilled water is provided to the computer room air conditioners through a centralized chilled water plant containing five chillers capable of providing 5,000 tons of chilled water to the facility. During our observation, the facility was providing chilled water solely to data center spaces, using two 1000-ton

chillers. The host was utilizing a 1000-ton York Millennium chiller installed in 2001 and a 1000-ton Trane unit, originally installed in the facility in 1986. The chillers are supported with a centralized six-cell condenser water system. The tower fans are equipped with variable speed drives.

The most common metric for comparing chilled water system efficiency is kW/ton of chilled water delivered to the data center. For this system, the electrical power in kW for the chillers, pumps, and fans involved in chilled water production and the chilled water flow to the data center were determined.

Energy intensity of the IT equipment in terms of Watts per square foot of raised floor was determined. At approximately 29 W/sq.ft., Facility 14 was slightly more intense than the average of other case studies to date. Figure 3 provides a summary of the intensities that were measured.




Figure 3 Computer Load Densities

In calculating energy intensity, the definitions used by the Uptime Institute are utilized. This excludes raised floor areas that are used for storage, control centers, etc.

Much of the load information was available through existing monitoring capability of the building controls systems and individual equipment read outs. In contrast to other case studies, the host site's staff was able to obtain most of the measured values.

Benchmarks are useful for several purposes:

- ◆ Providing a baseline to track performance over time

- ◆ Identifying the most energy intensive systems and components
- ◆ Uncovering operating and maintenance problems
- ◆ Finding energy and reliability operating and retrofit improvement opportunities
- ◆ Comparing performance to benchmarks observed elsewhere
- ◆ Determining energy intensity trends in computing equipment over time
- ◆ Establishing efficiency improvement goals based upon benchmark information



Figure 4 Computer Rack

- ◆ Establishing operating and design targets for future projects

The data center energy end use breakdown is shown below in figure 5.




Figure 5 Energy End Use Breakdown



Since most data centers are not fully occupied with Data processing equipment, it is difficult to determine how energy intensive they would be when full. To attempt to quantify the fully loaded condition, a qualitative estimate of occupancy of the floor space and the loading of the racks was made. In addition, the remaining electrical capacity from the UPS and Power Distribution Units were noted. This comparison was done for other case studies and the results are shown below in figure 6:




Figure 6 Projected computer load

Of primary interest in the facility was the HVAC system serving the data center areas, but other benchmarks were also obtained. The accuracy and completeness of data varies, based upon the accuracy of the various sources of power readings, measurement methods, access, and ease of measurement. Nonetheless, the data is sufficiently accurate to determine approximate energy intensities and end uses.

A number of general recommendations and observations for possible efficiency improvements or further study for this facility are provided in this report. The observations do not represent a comprehensive energy audit, but rather represent specific opportunities for improvement, for further study, or for use in future modifications or new construction.

II. DEFINITIONS

Data Center Facility	A facility that contains both central communications equipment, and data storage and processing equipment (servers) associated with a concentration of data cables. Can be used interchangeably with Server Farm Facility.
Server Farm Facility	A facility that contains both central communications equipment, and data storage and processing equipment associated with a concentration of data cables. Can be used interchangeably with Data Center Facility. Also defined as a common physical space on the Data Center Floor where server equipment is located (i.e. server farm).
Data Center Floor / Space	Total footprint area of controlled access space devoted to company/customer equipment. Includes aisle ways, caged space, cooling units, electrical panels, fire suppression equipment, and other support equipment. Per the Uptime Institute Definitions, this gross floor space is what is typically used by facility engineers in calculating a computer load density (W/sf). ¹
Data Center Occupancy	This is based on a qualitative estimate on how physically loaded the data centers are.
Data Center Cooling	Electrical power devoted to cooling equipment for the Data Center Floor space
Data Center Server/Computer Load	Electrical power devoted to equipment on the Data Center Floor. Typically the power measured upstream of power distribution units or panels. Includes servers, switches, routers, storage equipment, monitors, and other equipment.
Computer/Server Load Measured Energy Density	Ratio of actual measured Data Center Server Load in Watts (W) to the square foot area (ft ² or sf) of Data Center Floor. Includes vacant space in floor area.
Computer /Server Load Projected Energy Density	Ratio of forecasted Data Center Server Load in Watts (W) to square foot area (ft ² or sf) of the Data Center Floor if the Data Center Floor were fully occupied. The Data Center Server Load is inflated by the percentage of currently

¹ Users look at watts per square foot in a different way. With an entire room full of communication and computer equipment, they are not so much concerned with the power density associated with a specific footprint or floor tile, but with larger areas and perhaps even the entire room. Facilities engineers typically take the actual UPS power output consumed by computer hardware and communication equipment in the room being studied (but not including air handlers, lights, etc.) and divide it by the gross floor space in the room. The gross space of a room will typically include a lot of areas not consuming UPS power such as access aisles, white areas where no computer equipment is installed yet, and space for site infrastructure equipment like Power Distribution Units (PDU) and air handlers. The resulting gross watts per square foot (watt/ft²-gross) or gross watts per square meter (watt/m²-gross) will be significantly lower than the watts per footprint measured by a hardware manufacturer in a laboratory setting.

occupied space.

Cooling Load Tons

A unit used to measure the amount of cooling being done. Equivalent to 12,000 British Thermal Units (BTU) per hour.

Chiller Efficiency

The power used (kW), per ton of cooling produced by the chiller.

Cooling Load Density

The amount of cooling (tons) in a given area (ft² or sf).

III. Introduction

In order to gain a better understanding of the energy requirements associated with the increasing use and compaction of data processing equipment in data centers, the New York State Electric Research and Development Authority (NYSERDA) sponsored Lawrence Berkley National Laboratory (LBNL) to conduct two case studies in New York to determine actual energy use and to identify energy efficiency opportunity in the facilities. This study compares results to other case studies performed in New York and California. The goals of the studies included obtaining comparative benchmarks and eventually determining current best practices.

LBNL contracted with Syska Hennessy Group, a data center design firm with considerable data center design experience and conveniently located in New York City. Their role was to assist in collecting site data and obtaining measured electrical use for the facility. For this case study and benchmarking, a facility was selected where it was thought that most of the metrics of interest could be readily obtained through existing building management systems, or direct readout of equipment. Energy data was in large part provided directly by the facility engineering staff, then analyzed and compiled by LBNL and Syska. This is the first case study to facilitate “self-benchmarking” of a data center.

IV. Data center Overview

A large financial corporation with data center operations in Manhattan, NY volunteered to participate in the study. This facility is a ten story; steel frame, poured concrete building with a pre-cast curtain wall containing three floors of critical data center space with the remainder office space. The facility was built in 1986. The facility includes 750,000 square feet (sf). Of this total 83,471 ft² is raised floor for the support of data processing equipment. The remaining space is office, cafeteria, and other support space (equipment rooms, supply storage, etc.)

Some general information regarding the function and capacities of this facility are as follows:

- ◆ A central chilled water central plant provides cooling for the data center. Cooling for the office spaces is supplied from the same system as the data center so that isolating the cooling for the data center spaces only is very difficult.
- ◆ Two Redundant 69kV Electrical Feeds/facility draws approximately 2800 kW
- ◆ Two Independent On-site Substations for Site Service – ---MVA each providing six separate utility feeds using 15kV service.
- ◆ Ten 1400 KW Diesel Back Up Electrical Generators, with two additional generators under construction

- ◆ UPS Systems of 1500 KVA and 3000 KVA, total of six units
- ◆ Chiller Plant
 - Chiller (electric) capacity of 5,000 tons and
 - Six cell Cooling tower
 - Centralized chilled water plant consisting of 5 chillers, connected to a centralized condenser water system.
 - Chilled water pumped into common header for distribution throughout the entire facility.
- ◆ Energy Management Systems – Siemens Apogy BCS

The data center spaces contained various computing equipment with large concentrations of rack and cabinet mounted servers. Special cabinets with small auxiliary fans on top housed some of the servers. Some of the cabinets have the ability to control the fans based upon cabinet temperature. The cabinets have openings through the raised floor underneath the servers and allow a limited amount of cool air to enter through the front of the cabinet.

When predetermined temperatures are reached, thermostats activate the fan array mounted on the cabinet top to increase cool airflow through the cabinet. The end user felt this arrangement would allow airflow to equipment where heating was actually occurring as opposed to being routed to all equipment racks based on a simple mechanical airflow-balancing scheme. Not all racks were supplied with the auxiliary cooling apparatus, however, so further study would be required to determine its effectiveness.



Figure 7 Top of cabinet showing fans

It is believed that the auxiliary fan systems could conserve fan energy by allowing the facility to be more selective in CRAC unit utilization (i.e. some CRAC units may be able to be shut off). The presence of these auxiliary systems will not reduce chilled water use (use may actually increase as a result of additional fan energy), as they do not change the heat generated by the data processing equipment, nor do they change the environment in which the heat is rejected. The small cabinet fans do require some additional fan energy, which is provided through the UPS power to the IT equipment. The likely benefit for these systems is in reducing equipment hot spot temperatures through increased cabinet airflow.

V. ENERGY USE

WHOLE BUILDING ENERGY USE

The whole building electricity end use is shown in table 1 below. The whole building consumes an average of 8.5 MW of electricity. Here, the load associated with the data center includes the computer loads and infrastructure loads serving the data center spaces only. The remainder of the load (balance of building load) serves the rest of the building, consisting mostly of office space.

			Percent
TOTAL FACILITY kW	8500	kW	100%
TOTAL DATA CENTER kW	4785	kW	56%
BALANCE OF BUILDING LOAD kW	3715	kW	44%

TABLE 1 WHOLE BUILDING ENERGY

DATA CENTER ENERGY END USE

For the data center, the computer load (UPS load), HVAC fan energy, HVAC chilled water, lighting, UPS losses, and standby generator losses were determined. Figure 8 shows the relative breakdown of energy use in the data center.




Figure 8 Data Center 14 Energy Balance

	Measurement	Units	%
Computer Loads (UPS)	2434	kW	51%
UPS loss	212	kW	4%
HVAC – Chilled Water	1334	kW	28%
HVAC – Fan Energy	685	kW	14%
Lighting & Misc.	91	kW	2%
Standby Generator losses	30	kW	1%
Whole building power	8500	kW	100%
Whole building intensity	11.3	W/sf	

Table 2 Electrical Energy Breakdown

Assumptions

The following assumptions were made in arriving at the end use breakdown.

1. All air flows into the data centers were disregarded and assumed to be insignificant in meeting cooling requirements for data center spaces.
2. All plug loads were disregarded, as we did not identify any significant loads attached to non-PDU power sources.
3. All chilled water tonnage measured by the Siemens Control System was assumed routed to CRAC units, as detailed by the facility host.
4. The chilled water riser to the CRAC units was assumed to supply loads located in

the data centers, floor 6-8 only.

For this case study, the HVAC systems including chilled water, and fan energy accounted for approximately 42% of the electrical power associated with the data center. This performance is about average compared to other case studies to date. Lighting and plug loads are a very small percentage of the total energy, at 2% and this is consistent with most other data centers in the study. The UPS losses are 4% of the total data center power. The loss, however, is approximately 8% of the electrical power going to the computers, or 212 kW, or approximately 1800 MWh/yr. With an average electricity rate of approximately \$0.13/kW, this means that UPS losses amount to over \$2 million annually. Improving UPS efficiency would clearly result in significant savings.

The performance of the HVAC system can be evaluated based upon established metrics that can be used to directly compare system efficiency. Electrical load intensity for cooling is often represented by W/sf. However, another interesting metric for evaluating the efficiency of the data center cooling can be represented as a ratio of HVAC electrical power to the total power for the data center.

DATA CENTER 14 EFFICIENCY METRICS

Metric	Value	Units
Data Center HVAC Power	2019	kW
Total Data Center Power	4786	kW
Ratio of Cooling kW : Total Data Center kW	.42	--

Table 3 Efficiency Metrics

The overall data center computer load intensity is slightly above other measured intensities. The “cooling efficiency”, which is the efficiency normalized to the computer power is .42 (Cooling kW/total kW). This is in the top third of data centers in the study and should be investigated further.

For comparison purposes, the cooling efficiencies measured at other facilities are shown below. In this figure, the last entry represents this data center (Note there is no correlation to the facility numbers in the other figures – this is merely provided to show the range of values).




Figure 9 Cooling Efficiency at Other Data Centers

This metric may provide an indication of HVAC efficiency where lower ratios signify that more computing is provided for a given amount of cooling, however, the intensity and computing capability of the data processing equipment and other minor loads are also factors. UPS efficiencies drop off dramatically at part load conditions. A comparison of measured UPS efficiency at other case study sites is provided in figure 10.




Figure 10 Measured UPS efficiency as a function of Load Factor

Historical Data

The host facility staff routinely tracks total facility electric use as well as the W/sf. Recent load information was provided for the study. As expected, the total facility energy use is relatively constant, as business requirements are consistent and hardware changes are applied over a large inventory. The user consistently adds, changes or deletes equipment, but due to the large quantity of equipment and the size of the facility, trends in changes of electrical power consumption due to processing equipment changes would reveal themselves over longer periods of time, perhaps annually.

The host facility also manages its distribution of power consuming IT devices on the data center floors with the objective of minimizing space consumed by the equipment as well as balancing PDU loads on each floor. For planning purposes, the host does not use unoccupied spaces in the W/sf calculations, focused solely on the intensity of utilized areas. The business goals are to minimize the use of data center floor space, allowing additional equipment to be added, as business conditions require. The host's excess floor space would serve to lower the present actual W/sf calculation. The excess space is regarded as a corporate asset with some increased operating cost for CRAC unit fan energy.


VI. ENERGY EFFICIENCY OPPORTUNITIES:

General Recommendations

Efficient Chilled Water Systems

Water cooled chillers offer enormous energy savings over air cooled chillers. Since the chiller is being cooled by lower temperature media, it can reject heat more easily, and does not have to work as hard. Though the addition of a cooling tower adds maintenance costs associated with the water treatment, we have found that the energy savings outweigh the maintenance costs. Within the options of water cooled chillers, variable speed centrifugal are the most energy efficient, because they can operate very efficiently at low loads. The graph below compares the energy performance of various chiller types.

Comparison of Typical Chiller Efficiencies over Load Range



Chiller 1	250-Ton, Screw, Standard Efficiency, Air Cooled
Chiller 2	216 Ton, Screw, Water Cooled
Chiller 3	227-Ton, Centrifugal, Constant Speed, Water Cooled
Chiller 4	227-Ton, Centrifugal, Variable Speed, Water Cooled

Figure 11 Chiller Comparisons

Though there are efficient air cooled chillers, the larger size of water cooled chillers has resulted in more care given to efficiency and life cycle costs compared to air cooled chillers.

The selection of the auxiliary equipment, including the cooling tower, pumps, and pumping strategy should also be considered carefully. For example, variable speed fans on cooling towers allow for optimized cooling tower control. Premium efficiency motors and high efficiency pumps are recommended, and variable speed pumping is a ripe opportunity for pump savings. Variable pumping strategies can be achieved in a primary/secondary scheme, where the primary pumps operate at constant speed and directly feed water to the chiller, and the secondary pumps are variable speed and serve the air handling units. A more energy efficient scheme is primary-only variable speed pumping strategy. Pumping savings are based on the cube law: pump power is reduced by the cube of the reduction in pump speed, which is directly proportional to the amount of fluid pumped.

A primary only variable pumping strategy must include a bypass valve that ensures minimum flow to the chiller, and the use of two-way valves at the air handling units in order to achieve lower pumping speeds. The control speed of the bypass valve should also meet the chiller manufacturers recommendations of allowable turndown, such that optimum chiller efficiency is achieved. This basically means that the flow through the chiller should be varied slow enough such that the chiller is able to reach a quasi-steady state condition and able to perform to its maximum efficiency. The diagram below describes the primary-only variable speed pumping strategy.

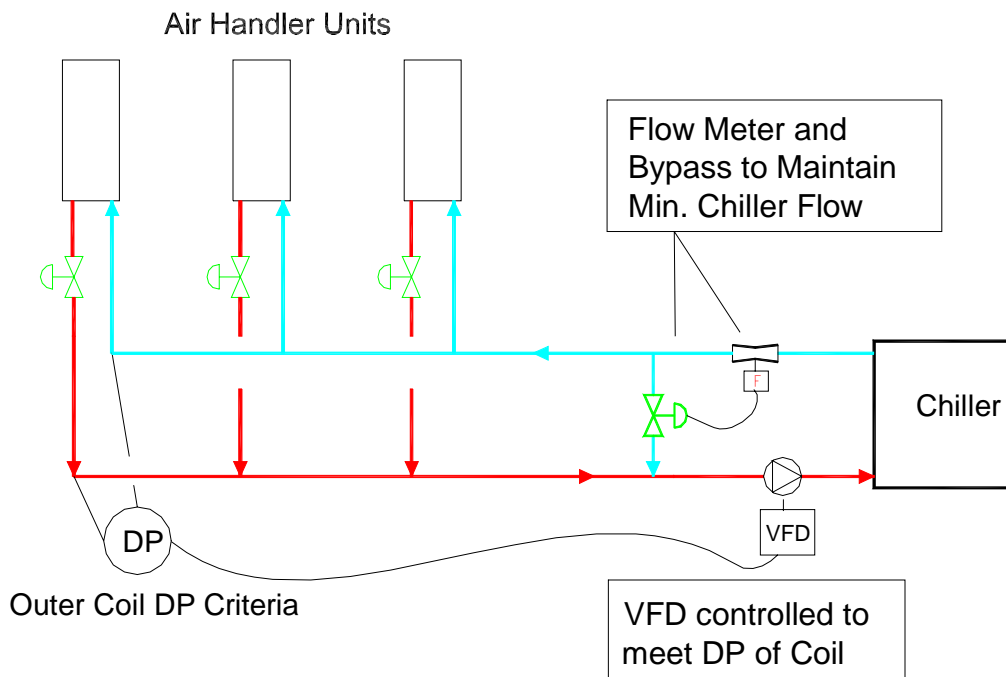


Figure 12 Air Handler Units

Air Management

A traditional method of cooling data centers employs an under floor system fed by CRAC units. There are a number of potential problems with such systems: an under floor system works on the basis of thermal stratification. This means that as the cool air is fed from the under floor, it absorbs energy from the space, warming up as a result, and rises. In order to take advantage of thermal stratification, the return air must be collected at the ceiling level. CRAC units often have low return air grills, and are therefore simply recirculating cool or moderately warmed air. Furthermore, they are often located along the perimeter of the building, and not dispersed throughout the floor area where they can more effectively treat warm air. One alternative is to install transfer grills from the ceiling to the return grill. near the ceiling.



Figure 13 This under floor area severely restricts air flow

Another common problem with under floor supply is that the under floor becomes congested with cabling, increasing the resistance to air flow. This results in an increase in fan energy use. A generous under floor depth is essential for effective air distribution. Figure 13 illustrates the potential for increased pressure gradients and

airflow restrictions that could cause local hot spot conditions and the need for increased fan energy.

An additional common problem in data centers is in placement of floor tiles with openings. Static and dynamic air pressures under raised floors are difficult to predict. Placement of openings close to CRAC units may have unintended consequences. Air flow modeling may help solve many placement problems.

An alternative to under floor air distribution is high velocity overhead supply, combined with ceiling height return. A central air handling system can be a very efficient air distribution unit. Design considerations include using VFDs on fans, low pressure-drop filters, and coils. An additional advantage of a central air handling system is that it can be specified with an economizer function. Many data centers can benefit from increased use of air side economizing, depending upon climate.

Humidity and Temperature Control

Another common problem identified with CRAC units is that they are often fighting each other in order to maintain a constant humidity setpoint. Not only is a constant humidity setpoint unnecessary for preventing static electricity (the lower limit is more important), but also it uses extra energy. A central air-handling unit has better ability to control overall humidity than distributed CRAC units.

Air Management – Rack Configuration

Server rack configuration also dictates air management strategies in data centers. It is more logical for the aisles to be arranged such that servers' backs are facing each other, and servers' fronts are facing each other. This way, cool air is drawn in through the front, and hot air blown out the back (assuming a front to back server). The Uptime Institute has published documents describing this method for air management.

Commissioning of New Systems and Optimized Control Strategies

Many times the predicted energy savings of new and retrofit projects are not fully realized. Often, this is due to poor and/or incomplete implementation of the energy efficiency recommendations. Commissioning is the process of ensuring that the building systems perform as they were intended to by the design. Effective commissioning actually begins at the design stage, such that the design strategy is critically reviewed. Either the design engineer can serve as the commissioning agent, or a third party commissioning agent can be hired. Commissioning differentiates from standard start-up testing in that it ensures systems function well relative to each other. In other words, it employs a systems approach.

Many of the problems identified in building systems are often associated with controls. A good controls scheme begins early in the design. In our experience, an effective controls design includes 1) a detailed points list, with accuracy levels, and sensor types, and 2) a detailed sequence of operations. Both of these components are essential for successfully implementing the recommended high efficiency chilled water system described above.

It is also possible that computer room air conditioners can be simultaneously cooling and humidifying – or heating and cooling at the same time. As noted below, however, it appears that cooling is not being provided by the CRAC units for the data center area examined.

Though use of commissioning is not uniformly adopted, various organizations have developed standards and guidelines. Such guidelines are available through organizations like Portland Energy Conservation Inc., at www.peci.org, or ASHRAE, Guideline 1-1996.

FACILITY 14 OBSERVATIONS

Computer Room Air Conditioning

Openings were observed sealed around many pieces of computing equipment through the raised floor utilizing a fire stop material. Sealing floor openings can improve efficiency by directing air through floor tiles to where it is needed. An air management scheme as described above directed the cool air and reduced the amount of air leaking from the raised floor. This lessened the fan energy required to provide sufficient air flow to the equipment cabinets.

Facility 14 did not utilize humidity control and data center spaces were maintained at approximately 68 deg F. Data center temperature and humidity control may be an opportunity for improvement. The host remarked he was considering raising the set

points. Studies have shown that the electrical components in data centers can withstand significantly higher temperatures.

In Facility 14, the return air was collected approximately six feet off of the raised floor. The CRAC units can more efficiently remove heat by collecting the heated air near the ceiling and taking advantage of the thermal stratification that occurs in the data center. A transfer grill or duct forcing the return to collect air near the ceiling could be investigated.

In a few areas, floor tiles were placed in inappropriate locations, allowing short-circuiting of conditioned air back to the CRAC units.

We did not investigate or observe the underfloor conditions or level of congestion under the raised floor in Facility 14.

Computer rack/HVAC interface

Facility 14 utilizes direct under floor air introduction into the cabinet spaces, via opening located at each cabinet bottom. Under floor airflow to the cabinets is adjustable as observed at several of the cabinets, by means of a variable opening floor tile. The host sets the opening in accordance to his own judgment and criteria. Warmed air was exhausted through the cabinet top. The host felt that this direct air introduction was most effective. Although the host felt that this scheme for air-cooling was very efficient in light of the business requirements and overall equipment configuration, other rack configurations may prove to be more efficient. The facility seems to be standardizing on the cabinet systems described above. Other arrangements utilizing hot and cold isles, or new emerging rack systems may hold promise for increased efficiency. Even though the fan energy for the small fans on the cabinets currently used may be small, there is a magnifying effect due to the numbers of fans, UPS losses, fan heat, and resulting increased HVAC load. This may be an area for further investigation.

Chilled Water System

Consider overall pumping energy for various combinations of chillers. Investigate use of free cooling and efficient operation of cooling towers. Determine which chillers are the most efficient and as increasing chiller loads are encountered, activate in the least efficient chiller last. Various resources are available to provide guidance for chilled water systems, such as Cooltools:

<http://www.hvacexchange.com/cooltools/coolhome.htm>)

Computer room lighting

Consider use of standard lighting controls such as timers or occupancy sensors in data center areas. Many areas are unoccupied for long periods of time and comparable savings to office areas can be obtained. Consider reduced lighting levels and/or eliminating lighting in certain areas, especially in times of peak demand charges. Many

telecom facilities hosting multiple customers are utilizing lighting controls to only illuminate a customer area when needed. This enhances energy savings and security. Savings for the direct cost of the lighting as well as the cost of removing the heat produced by the lighting will be realized.

Computing Equipment

Investigate ability to power down unused data processing equipment. Specifying IT equipment with improved power supplies, and improved idle state performance can dramatically improve the overall energy consumption in the data center.

Appendix A: Data Center References

ACEEE, and CECS. 2001. Funding prospectus for "Analysis of Data Centers and their implications for energy demand". Washington, DC, American Council for an Energy Efficient Economy (ACEEE); Center for Energy and Climate Solutions (CECS). July 2001.

The paper includes an overview of data centers; discusses energy use, energy choices, and energy efficiency in data centers; potential impacts of data centers; present and future regulatory issues; and business opportunities in energy services.

Aebischer, B., R. Frischknecht, C. Genoud, A. Huser, and F. Varone. 2002a. Energy- and Eco-Efficiency of Data Centres. A study commissioned by Département de l'intérieur, de l'agriculture et de l'environnement (DIAE) and Service cantonal de l'énergie (ScanE) of the Canton of Geneva, Geneva, November 15.

The study investigates strategies and technical approaches to fostering more energy-efficient and environmentally sound planning, building and operating of data centres. It also formulate recommendations on how to integrate the findings in the legal and regulatory framework in order to handle construction permits for large energy consumers and promote energy efficiency in the economic sectors. Seventeen recommendations grouped in four topics are derived from study conclusions: Transfer of the accord into an institutionalised legal and regulatory framework; Energy-efficiency policies for all large energy consumers; Preconditions, and prerequisites; Operational design of voluntary energy policies.

Aebischer, B., R. Frischknecht, C. Genoud, and F. Varone. 2002b. Energy Efficiency Indicator for High Electric-Load Buildings. The Case of Data Centres.

Proceedings of the IEECB 2002. 2nd International Conference on Improving Electricity Efficiency in Commercial Buildings. Nice, France.

Energy per unit of floor area is not an adequate indicator for energy efficiency in high electric-load buildings. For data centres we propose to use a two-stage coefficient of energy efficiency $CEE = C1 * c2$, where C1 is a measure of the efficiency of the central infrastructure and c2 a measure of the energy efficiency of the equipment.

Anonymous. 2001. Model Data Center Energy Design Meeting. Austin Energy, Austin, TX, Feb 12-13.

http://www.austinenergy.com/business/energy_design_meeting.htm

Anonymous. 2002a. 7 x 24 Update: Design & Construction - Issues and trends in mission critical infrastructure design, planning and maintenance.

<http://www.facilitiesnet.com/BOM/Jan02/jan02construction.shtml>. July 23, 2002.
<http://www.7x24exchange.org/>.

- Anonymous. 2002b. Continuous Availability Review (CAR). The Uptime Institute: Computersite Engineering, Inc. <http://www.upsite.com/csepages/csecar.html>. July 22, 2002.
- Anonymous. 2002c. End-to-End Reliability Begins with the User's Definition of Success. The Uptime Institute. <http://www.upsite.com/TUIpages/editorials/endtoend.html>. July 22, 2002.
- Anonymous. 2002d. Mechanical Systems Diagnostic Review (MSDR). The Uptime Institute: Computersite Engineering, Inc. <http://www.upsite.com/csepages/csemsdr.html>. July 22, 2002.
- Anonymous. 2002e. Site Infrastructure Operations Review (SIOR). The Uptime Institute: Computersite Engineering, Inc. <http://www.upsite.com/csepages/cseior.html>. July 22, 2002.
- Baer, D. B. Emerging Cooling Requirements & Systems in Telecommunications Spaces, Liebert Corporation.
During the last several years, power density trends, and consequently thermal density trends in telecommunications spaces have become topics of increasing interest. This paper identifies several of the underlying drivers of these trends, project possible outcomes, and assess the impact on cooling system design for these spaces.
- Beck, F. 2001. Energy Smart Data Centers: Applying Energy Efficient Design And Technology To The Digital Information Sector. *Renewable Energy Policy Project (REPP): Washington, DC*. (November 2001 REPP).
Both utilities and data center owners face challenges in meeting electricity demand loads with required levels of reliability. However, the bursting of the high-tech stock bubble in 2000 and the 2001 U.S. economic downturn has slowed expansion of data centers. This provides time and an opportunity to examine data center construction and operational practices with an eye toward reducing their energy demands through use of energy efficient technologies and energy smart design practices. As the economy recovers and the next data center rush approaches, best practices can reduce energy use while maintaining or even increasing data center reliability. Energy demands of data centers that support the digital information- and communications-based economy need not be as high as some predict. In fact, data center power demands could be reduced by 20 percent with minimal efficiency efforts, and by 50 percent with more aggressive efficiency measures.
- Blount, H. E., H. Naah, and E. S. Johnson. 2001. Data Center and Carrier Hotel Real Estate: Refuting the Overcapacity Myth. Lehman Brothers: TELECOMMUNICATIONS, New York, June 7, 2001. <http://www.lehman.com>
An exclusive study examining supply and demand trends for data center and carrier hotel real estate in North America. Lehman Brothers and Cushman &

Wakefield have completed the first in a regular series of proprietary studies on telecommunications real estate (TRE), including carrier hotels and data centers.

Bors, D. 2000. Data centers pose serious threat to energy supply. *Puget Sound Business Journal (Seattle)* - October 9, 2000.

<http://seattle.bizjournals.com/seattle/stories/2000/10/09/focus5.html>

To cope with increasing energy demand from data centers, the author discussed feasibilities of two possible approaches: 1) energy industry approach by looking at alternative energy supply; 2) construction industry approach by looking at data center energy efficiency. To get there, it is worth investigating four distinct components. (I) Co-generation of power. Presently, standby diesel generators are required to maintain the desired level of reliability at most data center sites, but their exhaust makes most of these generators unacceptable for long-term power generation. (II) Fuel cells offer the promise of very clean emissions and the reasonable possibility for use as standby power. (III) Increased efficiency in data center power distribution systems. There are two separate items that are major contributors to data center power distribution system inefficiencies. The first, power distribution units (PDUs), are available with optional internal transformers that use less energy than the present cadre of K-rated transformers. The second, uninterruptible power systems (UPSs), come in a range of efficiency ratings. If the use of high-efficiency PDUs and UPSs are combined, they offer the potential of a 6 percent saving. (IV) Increased efficiency in mechanical cooling systems. In order to ensure data center reliability, mechanical equipment is often selected as a large number of small, self-contained units, which offers opportunities to improve efficiencies. (V) Reductions in energy use by computer, network and storage equipment. Computer manufacturers can do their part by creating computers with greater computational power per watt. They have been doing this for years as a side effect of hardware improvements, and they can do even better if they make it a goal.

Brown, E., R. N. Elliott, and A. Shipley. 2001. Overview of Data Centers and Their Implications for Energy Demand. Washington, DC, American Council for an Energy Efficient Economy, Center for Energy & climate Solutions (CECS). September 2001. <http://www.aceee.org/pdfs/datacenter.pdf.pdf>

The white paper discusses data center industry boom and energy efficiency opportunities and incentives in internet data centers. Emerging in the late 1990's, data centers are locations of concentrated Internet traffic requiring a high-degree of power reliability and a large amount of power relative to their square footage. Typically, power needs range from 10-40MW per building, and buildings are typically built in clusters around nodes in the Internet fiber-optic backbone. During the development boom in 1999 and 2000, projects averaged 6-9 months from site acquisition to operation, and planned operational life was 36 months to refit. Even high energy-prices were dwarfed by net daily profits of 1-2 million dollars per day for these buildings during the boom, creating little incentive for efficient use of energy.

- Callsen, T. P. 2000. The Art of Estimating Loads. *Data Center* (Issue 2000.04).
This article discusses the typical Data Center layout. It includes floor plan analysis, HVAC requirements, and the electrical characteristics of the computer hardware typically found in a Data Center.
- Calwell, C., and T. Reeder. 2002. Power Supplies: A Hidden Opportunity for Energy Savings (An NRDC Report). Natural Resources Defense Council, San Francisco, CA, May 22, 2002. <http://www.nrdc.org>
The article discusses the efficiency of power supplies which perform current conversion and are located inside of the electronic product (internal) or outside of the product (external). The study finds that most external models, often referred to as "wall-packs" or "bricks," use a very energy inefficient design called the linear power supply, with measured energy efficiencies ranging from 20 to 75%; that most internal power supply models use somewhat more efficient designs called switching or switch-mode power supplies; and that internal power supplies have energy efficiencies ranging from 50 to 90%, with wide variations in power use among similar products. Most homes have 5 to 10 devices that use external power supplies, such as cordless phones and answering machines. Internal power supplies are more prevalent in devices that have greater power requirements, typically more than 15 watts. Such devices include computers, televisions, office copiers, and stereo components. The paper points out that power supply efficiency levels of 80 to 90% are readily achievable in most internal and external power supplies at modest incremental cost through improved integrated circuits and better designs.
- Cratty, W., and W. Allen. 2001. Very High Availability (99.9999%) Combined Heat and Power for Mission Critical Applications. *Cinintel 2001*: 12.
<http://www.surepowersystem.com>
- Elliot, N. 2001. Overview of Data Centers and their implications for energy demand. Washington, DC, American Council for an Energy Efficient Economy. Jan 2001, revised June 10, 2001.
- Feng, W., M. Warren, and E. Weigle. 2002. The Bladed Beowulf: A Cost-Effective Alternative to Traditional Beowulfs. *Cluster2002 Program*. <http://www-unix.mcs.anl.gov/cluster2002/schedule.html>; public.lanl.gov/feng/Bladed-Beowulf.pdf
Authors present a novel twist to the Beowulf cluster - the Bladed Beowulf. In contrast to traditional Beowulfs which typically use Intel or AMD processors, the Bladed Beowulf uses Transmeta processors in order to keep thermal power dissipation low and reliability and density high while still achieving comparable performance to Intel- and AMD-based clusters. Given the ever increasing complexity of traditional super-computers and Beowulf clusters; the issues of size, reliability, power consumption, and ease of administration and use will be "the" issues of this decade for high-performance computing. Bigger and faster machines are simply not good enough anymore. To illustrate, Authors present the

results of performance benchmarks on the Bladed Beowulf and introduce two performance metrics that contribute to the total cost of ownership (TCO) of a computing system - performance/power and performance/space.

Frith, C. 2002. Internet Data Centers and the Infrastructure Require Environmental Design, Controls, and Monitoring. *Journal of the IEST* **45**(2002 Annual Edition): 45-52.

Internet Data Centers and the Infrastructure Require Environmental Design, Controls, and Monitoring. The author points out that specifications and standards need to be developed to achieve high performance for mission-critical internet applications.

Gilleskie, R. J. 2002. The Impact of Power Quality in the Telecommunications Industry. Palm Springs, CA, June 4. <http://www.energy2002.ee.doe.gov/Facilities.htm>
The workshop addresses the unique issues and special considerations necessary for improving the energy efficiency and reliability of high-tech data centers. This presentation addresses impacts of power quality including voltage sags, harmonics, and high-frequency grounding in telecommunication industry.

Grahame, T., and D. Kathan. 2001. Internet Fuels Shocking Load Requests. *Electrical World* **Vol. 215** (3): 25-27.

http://www.platts.com/engineering/ew_back_issues.shtml

This article discusses the implications of the increase for power demand by the Internet's traffic growth on utility planning, operation, and financing.

Greenberg, D. 2001. Addendum to ER-01-15: A Primer on Harmonics. E-SOURCE, Boulder, Colorado, September 2001.

The electrical distribution systems of most commercial and industrial facilities were not designed to operate with an abundance of harmonics-producing loads. In fact, it is only within recent years that such loads have become widespread enough for industry to take notice and to begin to develop strategies to address the problems that harmonics can create. By 1992, concern about the issue had grown sufficiently that the Institute for Electrical and Electronic Engineers (IEEE) developed and published its standard 519, "IEEE Recommended Practices and Requirements for Harmonic Control in Electrical Power Systems," which established an approach for setting limits on the harmonic voltage distortion on the utility power system and on the harmonic currents created individual power consumers. Since that time, the electronic loads that give rise to harmonic currents have grown dramatically and are projected to continue growing for the foreseeable future. This being the case, there is and will continue to be a market for technological solutions to the problems that harmonics can cause.

Gross, P. 2002. Needed: New Metrics. *Energy User News*.

http://www.energyusernews.com/eun/cda/articleinformation/features/bnp__features__item/0,2584,82741,00.html

- Gruener, J. 2000. Building High-Performance Data Centers. *Dell Magazines - Dell Power Solutions* (Issue 3 "Building Your Internet Data Center").
http://www.dell.com/us/en/esg/topics/power_ps3q00_1_power.htm;
http://www.dell.com/us/en/esg/topics/power_ps3q00-giganet.htm
 The introduction of Microsoft SQL Server 2000 is a milestone in the race to build the next generation of Internet data centers. These new data centers are made up of tiers of servers, now commonly referred to as server farms, which generally are divided into client services servers (Web servers), application/business logic servers, and data servers supporting multiple instances of databases such as SQL Server 2000.
- Hellmann, M. 2002. Consultants Face Difficult New Questions in Evolving Data Center Design. *Energy User News*.
http://www.energyusernews.com/CDA/ArticleInformation/features/BNP__Features__Item/0,2584,70610,00.html
 While few data center design projects are alike, there are always the twin challenges of "power and fiber." And sometimes, even local politics and human factors. The paper suggested that the consultant should be brought in as soon as a business case is established so criteria can be established and a concept can be developed, priced, and compared to the business case. A planning is necessary before moving on to site selection and refine the concept and again test the business case.
- Howe, B., A. Mansoor, and A. Maitra. 2001. Power Quality Guidelines for Energy Efficient Device Application - Guidebook for California Energy Commission (CEC). Final Report to B. Banerjee, California Energy Commission (CEC).
 Energy efficiency and conservation are crucial for a balanced energy policy for the Nation in general and the State of California. Widespread adaptation of energy efficient technologies such as energy efficient motors, adjustable speed drives, improved lighting technologies will be the key in achieving self sufficiency and a balanced energy policy that takes into account both supply side and demand side measures. In order to achieve the full benefit of energy efficient technologies, these must be applied intelligently, and with clear recognition of the impacts some of these technologies may have on power quality and reliability. Any impediment to the application of these energy efficient technologies by the customers is not desirable for the overall benefit to energy users in California. With that in mind EPRI and CEC has worked to develop this guidebook to promote customer adaptation of energy efficient technologies by focusing on three distinct objectives. 1) Minimize any undesirable power quality impacts of energy-saving technologies; 2) Understand the energy savings potential of power quality-related technologies. These include: Surge Protective Devices (SPDs) or Transient Voltage Surge Suppressors (TVSS), Harmonic Filters, Power Factor Correction Capacitors, Electronic Soft Starters for Motors; and 3) How to evaluate "black box" technologies

Intel. 2002. Planning and Building a Data Center - Meeting the e-Business Challenge. Intel Corp.

http://www.intel.com/network/idc/doc_library/white_papers/data_center/. Aug 01, 2002.

The paper discusses the keys to success of Internet Service Providers (ISPs) that include 1) Achieve the economies of scale necessary to support a low price business model; 2) Offer added value, typically in the form of specialized services such as applications hosting to justify a premium price. This document provides a high-level overview of the requirements for successfully establishing and operating an Internet data center in today's marketplace. It offers some of the key steps that need to be taken, including project definition, prerequisites and planning. In order to construct a data center that can meet the challenges of the new market, there are three basic areas of data center definition and development: 1) Facilities: including building, security, power, air-conditioning and room for growth; 2) Internet connectivity: performance, availability and scalability; 3) Value-added services and the resources to support their delivery: service levels, technical skills and business processes. The aim is to provide customers with the physical environment, server hardware, network connectivity and technical skills necessary to keep Internet business up and running 24 hours a day, seven days a week. The ability to scale is essential, allowing businesses to upgrade easily by adding bandwidth or server capacity on demand.

Koplin, E. 2000. Finding Holes In The Data Center Envelope. *Engineered Systems* (September 2000).

http://www.esmagazine.com/CDA/ArticleInformation/features/BNP__Features__Item/0,2503,8720,00.html

The paper addresses importance of environmental control in data center facilities. Maintaining data center availability requires absolutely reliable infrastructure. A significant amount of this is devoted solely to maintaining stable environmental parameters. And only constant, thorough regulation and testing of these parameters ensures the integrity of the data center "envelope."

Mandel, S. 2001. Rooms that consume - Internet hotels and other data centers inhale electricity. *Electric Perspectives* **Vol. 26** (No.3).

http://www.eei.org/ep/editorial/Apr_01/0401ROOM.htm

The article estimated that the amount of this data center space in the United States nearly doubled in 2000, totaling between 19 million and 25 million square feet by year-end, according to investment analysts. They say they expect another 10 million to 20 million square feet of new space to be added in 2001. Developers are asking electric utilities to supply the buildings with 100-200 watts of electricity per square foot. Since these data centers are new to the economy, there is little historical data on which to base estimates of electricity use for a facility. In addition, the dot.com world makes it difficult for the developer to say confidently how much electricity one of these internet hotels will use. Source One estimates that tens of billions of dollars worth of electric infrastructure improvements will be needed for data centers over the next few years and that

they will consume billions of dollars more worth of electricity. The energy costs are as high or higher than the actual lease costs. Indeed, 50-60 percent of the cost of building a data center is for the power, including batteries, backup generators, and air-conditioning, as well as the cost for utility construction.

Mitchell-Jackson, J. 2001. Energy Needs in an Internet Economy: A Closer Look at Data Centers, July, 2001.

This study explains why most estimates of power used by data centers are significantly too high, and gives measured power use data for five such facilities. Total power use for the computer room area of these data centers is no more than 40 W/square foot, including all auxiliary power use and cooling energy. There are two draft journal articles from this work, one focusing on the detailed power use of the data center we've examined in most detail, and the other presenting the aggregate electricity use associated with hosting-type data centers in the U.S.

Mitchell-Jackson, J., J. G. Koomey, B. Nordman, and M. Blazek. 2001. Data Center Power Requirements: Measurements From Silicon Valley. *Energy—the International Journal* (Under review).

<http://enduse.lbl.gov/Projects/InfoTech.html>

Current estimates of data center power requirements are greatly overstated because they are based on criteria that incorporate oversized, redundant systems, and several safety factors. Furthermore, most estimates assume that data centers are filled to capacity. For the most part, these numbers are unsubstantiated. Although there are many estimates of the amount of electricity consumed by data centers, until this study, there were no publicly available measurements of power use. This paper examines some of the reasons why power requirements at data centers are overstated and adds actual measurements and the analysis of real-world data to the debate over how much energy these facilities use.

Patel, C. D., C. E. Bash, C. Belady, L. Stahl, and D. Sullivan. 2001. Computational Fluid Dynamics Modeling of High Compute Density Data Centers to Assure System Inlet Air Specifications. Reprinted from the proceedings of the Pacific Rim ASME International Electronic Packaging Technical Conference and Exhibition (IPACK 2001), © 2001, ASME.

Due to high heat loads, designing the air conditioning system in a data center using simple energy balance is no longer adequate. Data center design cannot rely on intuitive design of air distribution. It is necessary to model the airflow and temperature distribution in a data center. This paper presents a computational fluid dynamics model of a prototype data center to make the case for such modeling.

Patel, C. D., R. Sharma, C. E. Bash, and A. Beitelmal. 2002. Thermal Considerations in Cooling Large Scale High Compute Density Data Centers. 8th IThERM Conference. San Diego CA.

A high compute density data center of today is characterized as one consisting of thousands of racks each with multiple computing units. The computing units

include multiple microprocessors, each dissipating approximately 250 W of power. The heat dissipation from a rack containing such computing units exceeds 10 KW. Today's data center, with 1000 racks, over 30,000 square feet, requires 10 MW of power for the computing infrastructure. A 100,000 square foot data center of tomorrow will require 50 MW of power for the computing infrastructure. Energy required to dissipate this heat will be an additional 20 MW. A hundred thousand square foot planetary scale data center, with five thousand 10 KW racks, would cost ~\$44 million per year (@ \$100/MWh) just to power the servers & \$18 million per year to power the cooling infrastructure for the data center. Cooling design considerations by virtue of proper layout of racks can yield substantial savings in energy. This paper shows an overview of a data center cooling design and presents the results of a case study where layout change was made by virtue of numerical modeling to avail efficient use of air conditioning resources.

PG&E. 2001. Data Center Energy Characterization Study. Pacific Gas and Electric Company (subcontractor: Rumsey Engineers), San Francisco, Feb. 2001. Rumsey Engineers, Inc. and PG&E have teamed up to conduct an energy study as part of PG&E's Data Center Energy Characterization Study. This study will allow PG&E and designers to make better decisions about the design and construction of data centers in the near future. Three data centers in the PG&E service territory have been analyzed during December 2000 and January 2001, with the particular aim of determining the end-use of electricity. The electricity use at each facility was monitored for a week each. At the end of the report are a set of definitions, which explain the terms used and the components in making each calculation. The three data centers provide co-location service, which is an unmanaged service that provides rack space and network connectivity via a high capacity backbone. About half or more of the electricity goes to powering the data center floor, and 25 to 34 percent of the electricity goes to the heating, air conditioning and ventilation equipment. The HVAC equipment uses a significant amount of power and is where energy efficiency improvements can be made. All three facilities use computer room air conditioning (CRAC) units, which are stand-alone units that create their own refrigeration and circulate air. A central, water-cooled chilled water system with air handlers and economizers can provide similar services with roughly a 50% reduction in cooling energy consumption. Energy density of the three buildings had an average of 35 W/sf. The cooling equipment energy density for the data center floor alone averaged at 17 W/sf for the three facilities. The average designed energy density of the three data centers' server loads was 63 W/sf, while the measured energy density was 34 W/sf. An extrapolated value was also calculated to determine what the server load energy density would be when fully occupied. The average extrapolated energy density was 45 W/sf. Air movement efficiency varies from 23 to 64 percent between the three facilities. Cooling load density varies from 9 to 70 percent between the three facilities.

Planet-TECH. 2002. Technical and Market Assessment for Premium Power in Haverhill. Planet-TECH Associates for The Massachusetts Technology Collaborative,

www.mtpc.org, Westborough, MA 01581-3340, Revision: February 20, 2002.
http://www.mtpc.org/cluster/Haverhill_Report.pdf ; <http://www.planet-tech.com/content.htm?cid=2445>

This study is pursued under contract to the Massachusetts Technology Collaborative, in response to a request for a "Technical and Market Assessment". It seeks to determine if the provisioning of "premium power" suitable for data-intensive industries will improve the marketability of a Historic District mill building in Haverhill. It is concluded that such provisioning does improve the marketability, however, not to a degree that is viable at this time. Other avenues for energy innovation are considered and recommendations for next steps are made.

RMI, and DR International. 2002. Energy Efficient Data Centers - A Rocky Mountain Institute Design Charrette. Organized, Hosted and Facilitated by Rocky Mountain Institute, with D&R International, Ltd. and Friends. Hayes Mansion Conference Center, San Jose, California.

Rapid growth of "mission critical" server-farm and fiber-optic-node data centers has presented energy service providers with urgent issues. Resulting costs have broad financial and societal implications. While recent economic trends have severely curtailed projected growth, the underlying business remains vital. The current slowdown allows us all some breathing room—an excellent opportunity to step back and carefully evaluate designs in preparation for surviving the slowdown and for the resumption of explosive growth. Future data center development will not occur in the first-to-market, damn-the-cost environment of 1999-2000. Rather, the business will be more cost-competitive, and designs that can deliver major savings in both capital cost (correct sizing) and operating cost (high efficiency)—for both new build and retrofit—will provide their owners and operators with an essential competitive advantage.

Robertson, C., and J. Romm. 2002. Data Centers, Power, and Pollution Prevention - Design for Business and Environmental Advantage. The Center for Energy and Climate Solutions; A Division of The Global Environment and Technology Foundation, June 2002. <http://www.cool-companies.org>; <http://www.getf.org>
Computers and other electronic equipment will crash at the slightest disruption or fluctuation in their supply of electricity. The power system was not designed for these sensitive electronic loads and is inherently unable to meet the technical requirements of the information economy. For data centers, which play a central role in the information economy, crashing computers cause potentially catastrophic financial losses. The same voltage sag that causes the lights to dim briefly can cause a data center to go off-line, losing large sums of money, for many hours. Data center owners and their power providers must therefore solve several related technical and economic electric power problems. These are: 1) How to assure high-availability (24x7) power supply with a very low probability of failure; 2) How to assure practically perfect power quality; and 3) How to manage risk while minimizing capital and operating expenses

Roth, K. W., Fred Goldstein, and J. Kleinman. 2002. Energy consumption by office and telecommunications equipment in commercial buildings, Volume I: Energy Consumption Baseline. Arthur D. Little (ADL), Inc., 72895-00, Cambridge, MA, January 2002.

ADL carried out a "bottom-up" study to quantify the annual electricity consumption (AEC) of more than thirty (30) types of non-residential office and telecommunications equipment. A preliminary AEC estimate for all equipment types identified eight key equipment categories that received significantly more detailed studied and accounted for almost 90% of the total preliminary AEC. The Key Equipment Categories include: *Computer Monitors and Displays, Personal Computers, Server Computers, Copy Machines, Computer Network Equipment, Telephone Network Equipment, Printers, Uninterruptible Power Supplies (UPSs)*. The literature review did not uncover any prior comprehensive studies of telephone network electricity consumption or uninterruptible power supply (UPS) electricity consumption. The AEC analyses found that the office and telecommunications equipment consumed 97-TWh of electricity in 2000. The report concludes that commercial sector office equipment electricity use in the U.S. is about 3% of all electric power use. The ADL work also creates scenarios of future electricity use for office equipment, including the energy used by telecommunications equipment.

Sullivan, R. F. 2002. Alternating Cold and Hot Aisles Provides More Reliable Cooling for Server Farms. The Uptime Institute.

<http://www.uptimeinstitute.org/tuiaisles.html>


The creation of "server farms" comprising hundreds of individual file servers has become quite commonplace in the new e-commerce economy, while other businesses spawn farms by moving equipment previously in closets or under desktops into a centralized data center environment. However, many of these farms are hastily planned and implemented as the needed equipment must be quickly installed on a rush schedule. The typical result is a somewhat haphazard layout on the raised floor that can have disastrous consequences due to environmental temperature disparities. Unfortunately, this lack of floor-layout planning is not apparent until after serious reliability problems have already occurred.

The Uptime Institute. 2000. Heat-Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment. The Uptime Institute, Version 1.0., <http://www.upsite.com/>. <http://www.uptimeinstitute.org/heatdensity.html>

This white paper provides data and best available insights regarding historical and projected trends in power consumption and the resulting heat dissipation in computer and data processing systems (servers and workstations), storage systems (DASD and tape), and central office-type telecommunications equipment. The topics address the special needs of Information Technology professionals, technology space and data center owners, facilities planners, architects, and engineers.

- Thompson, C. S. 2002. Integrated Data Center Design in the New Millennium. *Energy User News*.
http://www.energyusernews.com/CDA/ArticleInformation/features/BNP__Features__Item/0,2584,70578,00.html
Data center design requires planning ahead and estimating future electrical needs. Designers must accurately predict space and energy requirements, plus cooling needs for new generations of equipment. Importance of data center reliability is discussed.
- Wood, L. 2002. Cutting Edge Server Farms - The blade server debate.
[newarchitectmag.com](http://www.newarchitectmag.com).
<http://www.newarchitectmag.com/documents/s=2412/na0702f/index.html>. July 23, 2002. A blade is the industry term for a server that fits on a single circuit board, including CPU, memory, and perhaps a local hard disk. Multiple blades are plugged into a chassis, where each blade shares a common power supply, cooling system, and communications back plane. Multiple chassis can then be stacked into racks. By comparison, the conventional approach for rack-mounted servers involves only one server per chassis. A chassis cannot be smaller than one vertical rack unit (1U, or about 1.75 inches high). This limits you to 42 to 48 servers in a standard seven-foot rack. A typical blade chassis is much higher than 1U, but several can still be stacked in a rack, allowing upwards of 300 servers per rack, depending on the vendor and configuration. This compact design offers compelling advantages to anyone operating a high-density server farm where space is at a premium. Indeed, blades are the "next big thing" in servers, and it's probable that any given administrator will have to decide whether to adopt them in the near future.

**Appendix XIII - “High-Performance Data Centers – a Research
Roadmap”, Lawrence Berkeley National Laboratory Report No.
53483**



HIGH-PERFORMANCE DATA CENTERS

A RESEARCH ROADMAP

Developed by:

Lawrence Berkeley National Laboratory with input from industry partners representing data center facility design and operation firms, industry associations, research organizations, energy consultants, and suppliers to data centers

**William Tschudi, Tengfang Xu, Dale Sartor
And
Jay Stein – E Source**

Sponsored by:

**The California Energy Commission through the
Public Interest Energy Research (PIER) Program**



DISCLAIMER

The Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720 is a National Laboratory of the U.S. Department of Energy (DOE) managed by the Regents of the University of California for the U.S. Department of Energy under Contract Number DE-AC03-76SF00098. This report was prepared as an account of work sponsored by the California Energy Commission (Commission, Energy Commission) and pursuant to an M&O Contract with the United States Department of Energy (DOE). Neither the Regents of the University of California, nor the DOE nor the California Energy Commission nor any of their employees, contractors, or subcontractors, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by The Regents of the University of California, or the DOE, or the California Energy Commission. The views and opinions of the authors expressed herein do not necessarily state or reflect those of The Regents of the University of California, the DOE, or the California Energy Commission, or any of their employees, or the United States Government, or any agency thereof, or the State of California. This report has not been approved or disapproved by the Regents of the University of California, the DOE, or the California Energy Commission, nor has The Regents of the University of California, the DOE, or the California Energy Commission passed upon the accuracy or adequacy of the information in this report.



ROADMAP FOR PUBLIC INTEREST RESEARCH

For

HIGH-PERFORMANCE DATA CENTERS

William Tschudi, Tengfang Xu, Dale Sartor, and

Jay Stein – E Source

Building Technologies Department
Environmental Energy Technologies Division
Ernest Orlando Lawrence Berkeley National Laboratory
University of California
1 Cyclotron Road
Berkeley, California 94720-8134 USA

Table of Contents

EXECUTIVE SUMMARY	1
INTRODUCTION AND BACKGROUND	4
Data Center Definition	4
Case Studies and Prior Investigations	5
Workshops to Develop the Roadmap	5
RMI Data Center Charrette	6
Development of the Roadmap	7
Organization of this Roadmap	8
ROADMAP	8
Collecting, Analyzing, and Applying Data Center Market Information	8
Understanding the Data Center Market	8
Benchmarking Energy Use	9
Identifying and Developing Best Practices	15
Improving Facility Infrastructure Efficiency	16
Developing Better Monitoring and Control Tools	19
Electrical System Issues	20
HVAC System Issues	25
Making Better Use of Existing Energy Efficiency Guidance	30
Improving the Interface Between Building Systems and IT Equipment	30
Improving Efficiency of IT Equipment	34
Appendix A	38
References	40

Data Center Energy Research and Deployment Roadmap

Executive Summary

When California's electric utilities began receiving requests for huge electrical demands for data center facilities, it became evident that little information existed to validate actual data center electrical performance, or to see how the energy performance could be improved. As a result, California utilities and the California Energy Commission became interested in learning more about the data center market. Utility case studies and preliminary investigations confirmed that research with the objective of reducing the large, continuous electrical loads in data centers was clearly merited, however the role of public interest research for these types of facilities was not clear.

To tackle this problem, the Public Interest Energy Research (PIER) Industrial Program set out to define and prioritize energy efficiency research areas by engaging Data Center Industry professionals. In preparation of this roadmap, researchers from Lawrence Berkeley National Lab (LBNL) facilitated workshops, participated in industry forums, and researched energy issues related to data centers. As the topics in the roadmap were developed, opportunities for California public interest research and market transformation activities were the primary focus. Other research and standardization activities by others were noted, and it will be important to keep abreast of their progress as the California research agenda is advanced. In addition, data center professionals identified other parts of the energy efficiency puzzle that must be solved by the industry itself due to the highly specialized nature of much of the equipment in data centers. Even though the research in these areas will proceed through industry efforts, public interest encouragement may accelerate the development and adoption of new innovations.

There are many types of data centers in California ranging from corporate data centers in a wide range of industries, banks, telecommunication facilities, and Internet hosting facilities. Data centers are also found in other institutions such as research organizations, universities, National Laboratories, and government facilities. The crosscutting nature of the market involves many California industries that directly or indirectly rely on data centers, as well as the suppliers of goods and services for data centers.

The data center market when defined broadly represents a significant and important component of the California economy. Regardless of their mission and make-up, most data centers are much more energy



Figure 1 The National Energy Research Scientific Computing Center (NERSC) operated by LBNL

intensive than other buildings due to the high power requirements of the computing equipment and the infrastructure needed to support the computing equipment. Based on their energy density, large data centers more closely resemble industrial facilities than commercial buildings.

The roadmap development identified many areas where significant efficiency gains could be achieved through adoption of current best practices, better application of existing technology, and research into new technological solutions. The roadmap organizes these areas as follows:

1. Activities aimed at understanding the Data Center Market – The size and growth rate of the market as well as local concentrations of data centers is of interest to planners and implementers of electrical power generation and distribution.
2. The benefits of obtaining energy benchmarks – By monitoring and comparing the energy consumption of a variety of data centers, operators and designers will be able to learn what is possible to achieve.
3. Identification and promotion of best practices – Adopting current best practices in existing or new data centers will provide significant improvement in the short term.
4. Improving data center facility systems' efficiency - Facility systems containing both conventional equipment such as chillers, and specialty equipment such as Uninterruptible Power Supplies are far from optimal.
5. Improving the interface between building systems and IT Equipment-The systems that house and support electronic equipment in data centers are typically not designed to optimize the efficiency of the building infrastructure systems they interface with.
6. Improving the efficiency of IT Equipment - Energy use in data centers is dominated by the servers, hard drives, routers, and switches that are used to process, store, and transmit data. Efficiency improvements in IT equipment are compounded by secondary effects in HVAC and power supply facility systems

During the course of the roadmap development, data center experts suggested several promising research topics that were outside the scope of this roadmap, which is focused exclusively on improving the energy efficiency of data centers. Most of these topics, such as thermal storage for peak demand reduction and distributed generation, have the potential to provide other societal benefits, and as such, are under investigation by other parties outside of the data center industry so they were not included in the roadmap. Furthermore, it was recognized that research efforts were underway by various organizations and industry associations, such as iTherm, CEETHERM, (a collaboration between the University of Maryland and Georgia Tech) and major electronics companies. In some cases the roadmap cites these efforts, observing that there are other research efforts outside those included in this roadmap that will make important contributions toward solving the overall problem. Collaboration and awareness of developments by others will be important to make sure that the research undertaken in California is headed in the right direction.

Lastly, while developing the roadmap, we uncovered an important difference of opinion within the data center industry. The electronic equipment used by this industry is continually evolving. Some industry observers have noted that the energy intensity exhibited by this equipment (measured in Watts per square foot) is increasing. Others, noting the recent availability of more efficient microprocessors, have proposed that at some point in the future, the trend towards increasing energy intensity will either slow down, level off, or decline. In other cases, more powerful computing equipment, although itself more energy intensive, has replaced many other pieces of equipment resulting in a net decrease in energy use. We did not take a position in the debate over whether intensities will rise or fall. However, given such uncertainty regarding future electrical and cooling demands, our efforts were directed at identifying strategies that would allow for efficient data center operation regardless of how technology evolution and business conditions play out.

Introduction and Background

Data Center Definition

The market addressed by this roadmap employs a broad definition of the term "data center". Generally, we use the term data center to be a facility that contains concentrated equipment to perform one or more of the following functions: Store, manage, process, and exchange digital data and information. Such digital data and information is typically applied in one of two ways:

- ◆ Support the informational needs of large institutions, such as corporations and educational institutions.
- ◆ Provide application services or management for various types of data processing, such as web hosting, Internet, intranet, telecommunication, and information technology.

We do not consider spaces that primarily house office computers, including individual workstations, servers associated with workstations, or small server rooms, to be data centers. Generally, the data centers we include are designed to accommodate the unique needs of energy intensive computing equipment along with specially designed infrastructure to accommodate high electrical power consumption, redundant supporting equipment, and the heat dissipated in the process.

To accommodate these needs, data centers typically exhibit these characteristics:

- ◆ Physically house various types of IT equipment, such as computers, servers (e.g., web servers, application servers, database servers), main frame computers, switches, routers, data storage devices, load balancers, wire cages or closets, vaults, racks, and related equipment.
- ◆ Exhibit critical requirements for security and reliability.
- ◆ Most, but not all, data centers utilize raised floors or other specialized computer room air conditioning systems.
- ◆ Provide for redundant and uninterruptible power.

Other terms used to refer to facilities that meet the definition of data center as used in this roadmap include: computer center, data storage and hosting facility, server farm, data farm, data warehouse, co-location facility, co-located server hosting facility (CoLo), corporate data center, managed data centers, internet hotel, internet service provider (ISP), application service provider (ASP), full service provider (FSP), wireless application service provider (WASP), telecommunication hotel (or telco hotel), carrier hotel, internet hotel, telecommunications carriers, or other data networks.



Figure 2 Supercomputers
in LBNL's NERSC Center

Case Studies and Prior Investigations

A number of case studies have been performed to characterize the energy end use in various types of data centers¹. A large variation in energy intensity and energy efficiency of key systems was observed in the various facilities that were studied. Design features of the better performing systems were noted yet every facility had the potential for energy efficiency improvement. Recommendations for efficiency improvements were provided as part of the case studies. The recommendations and findings often identified common issues. Many of the issues noted in the

case studies suggested areas where further research could lead to much better performance.




Figure 3 Typical Data Center Energy End Use

Workshops to Develop the Roadmap

There are many stakeholders interested in data center energy consumption. These stakeholders include data center industry associations, research organizations, public interest groups (i.e. utilities, federal and state agencies), as well as the individual firms that operate data centers, and manufacture IT or data center specialized equipment.

Input was obtained from data center professionals and experts in order to develop this research roadmap for high-performance, energy-efficient data centers, and to validate research issues and possible actions identified through case studies. Their input was obtained throughout the project by conducting workshops, attending data center conferences and meetings, and interviewing data center professionals. Leading data center designers, specialty equipment suppliers, computer manufacturers, energy consultants, and industry associations were contacted to solicit input concerning the state of the data center industry, and for help in defining where public interest research could make a difference. Industry associations such as the 7 X 24 Exchange Organization (www.7x24exchange.org) and the Uptime Institute (www.uptime.com) participated and provided valuable input including research topics and possible actions to address them. These professionals also helped to prioritize the issues and possible actions.

Generally, two types of recommendations emerged:

◆ Recommendations to Improve performance through better use of existing technologies

Through case studies, a wide variation in energy performance using today's technologies was observed. As a result, the identification of current best practices and efforts to

¹ Case studies and summary information are available on LBNL's website: <http://datacenters.lbl.gov>.

influence the market to adopt best practices should be a high priority. Some roadmap recommendations involve better use of existing strategies and technologies.

◆ Research and Development for new solutions

To advance beyond the current best practices, research and development will be needed in a number of areas. Industry participants identified many areas and issues where new solutions are needed, both for improving efficiency and for handling expected increases in heat intensity.

RMI Data Center Charrette

The California Energy Commission, along with other public interest organizations, sponsored a data center “charrette²”. This event, hosted and coordinated by the Rocky Mountain Institute (RMI- www.rmi.org), assembled over 75 leading stakeholders in the data center market for three days of brainstorming and sharing their expertise and information. The charrette participants identified many of the same issues previously identified during the case study work and other workshops, but also introduced additional issues. Many ideas for needed research were presented from the chip level up through the facility power source. Some research topics are more appropriate for various areas of the data center industry to take the lead (e.g. Efficiency of “blade servers”), and others where public interest research is needed and appropriate (e.g. benchmarking energy use). The public interest research topics along with some topics for possible joint participation are included in this roadmap. RMI will report charrette findings separately.




Figure 4 RMI Charrette Announcement

² The term charrette describes a process widely used by architects to critique a design and brainstorm new solutions. Normally the charrette occurs early enough in the design process to allow improvements to be incorporated.

Development of the Roadmap

In addition to the workshops and design charrette described above, an extensive literature review further helped to identify trends, current practice, and also suggested areas where improvement is needed. References that were reviewed were annotated and included as appendix A. Through these activities, a list of issues with a bearing on energy efficiency was compiled and forms the basis of this roadmap. For each of the issues one or more suggested research and/or market transformation actions were developed. Some of the actions are intended to provide near term improvement by determining current best practices, creating new ways to use existing technology, overcoming barriers, and helping the market adopt energy efficient concepts. Other actions are longer term but have the potential to bring further dramatic efficiency improvement to the market. The participants at the RMI data center charrette (RMI 2003), for example, felt that an order of magnitude energy reduction was possible.

To achieve this level of improvement, it is likely that all elements in the data center from chip level through building systems, to building shell would need to be optimized. This level of improvement would require simultaneous and coordinated RD&D efforts involving all energy using devices and systems – for example; improving the efficiency of computer chips, computer power supplies, heat transfer through cabinets, HVAC systems, UPS systems, standby power reduction, more efficient computer code, etc. This lofty goal is unlikely to occur without strategic guidance given the fragmented nature of market, the number and variety of data center suppliers, and the evolving nature of the wide assortment of IT equipment. In addition there are numerous barriers to change. Issues such as fast track design and construction schedules, reliability at any cost, inertia to maintain proven (although inefficient) design, etc. are preventing advances in this market. However, large efficiency gains are possible in the areas identified in this roadmap. An integrated strategy that leverages public interest funding has the potential to achieve a dramatic efficiency gain.

Current understanding of data center energy efficiency in the industries and institutions that rely on them is very limited. Typically, data center professionals have a thorough understanding of issues related to power quality, reliability, and availability, but energy efficiency is not a high priority. The general lack of benchmark information, various definitions of energy intensity, together with traditional barriers limiting efficiency improvements in data centers immediately suggests areas for further research, development and market transformation. In addition, case study recommendations and other industry input point to many areas where large energy savings are possible.

Organization of this Roadmap

The topic descriptions in the roadmap are organized into the following categories:

- ◆ Collecting, analyzing, and applying data center market data
- ◆ Improving Facility Infrastructure Efficiency
- ◆ Improving Building Systems and IT Equipment Interfaces
- ◆ Improving IT Equipment Efficiency

For each roadmap activity, industry participants attempted to identify the activities most suited for public interest research, development, and demonstration (RD&D) actions recognizing that some areas in need of research are not appropriate for public interest efforts. Research and advancement in some areas can best be accomplished by industry efforts, such as improving idle state performance of computing equipment. Other activities are good candidates for public interest involvement since they would not otherwise be accomplished such as benchmarking performance across the various industries that operate data centers. Still other actions may be better accomplished through collaboration or setting standards of performance with data center IT or equipment suppliers. Examples of this category are potentially to develop more efficient power supplies for IT equipment or to develop more efficient specialty infrastructure systems.

Roadmap

Collecting, Analyzing, and Applying Data Center Market Information

Understanding the Data Center Market

Data centers house computers and IT equipment to provide functions such as data information storage, data processing, and information dissemination. Many definitions have been used to describe “data centers” [ACEEE and CECS 2001; Aebischer et al. 2002b; Blount et al. 2001; Brown et al. 2001; Callsen 2000; Elliot 2001; Gruener 2000; Intel 2002; Mitchell-Jackson 2001; Planet-TECH 2002; Robertson and Romm 2002]. With the boom associated with the Internet in the late ‘90’s came new names for data centers including “server farms”, “collocation facilities” and “telecommunication hotels”. Although data centers have been important to industries, institutions, and governmental agencies for some time, it was the Internet and mission-critical facilities that brought energy consumption in data centers to the forefront.

Electric power requirements for data centers became an important issue for three very different reasons. First, computer technology, primarily chip technology, was creating higher heat density in smaller and smaller geometries. The simultaneous compaction and increase in electrical power caused concern over the ability to cool future generations of IT equipment. Secondly, the facilities that support the Internet were requesting unrealistic levels of electrical power from utilities. That requested power, if it materialized, would have required major changes in electrical utility generation and distribution infrastructure. Third, IT professionals, data center

operators, and facility designers aggravated both situations by predicting huge increases in electrical demand for future computing equipment. Limited energy benchmarks in operating data centers confirm that present day energy use is much lower than predicted. When those high-expected loads did not materialize, over sizing of data center infrastructure resulted in inefficient operation in many data centers. If criteria are not developed to improve the understanding of near and longer-term electric load requirements, such inefficient operation is likely to continue into the future.

To first come to grips with the extent of this problem, the place to begin is to characterize the stock of data centers and their load intensity in California. These characteristics have turned out to be difficult to estimate. The market is characterized by constant change and there is no reliable source of market data covering all of the various types of data centers. Load intensity for data centers supporting the Internet fluctuates greatly with the rise and decline of dot COM companies but data center load intensity is also affected by the trends in computing capability and energy intensity within IT equipment. One scenario suggests that the total computing electrical load is increasing at a modest pace and being compacted into a smaller number of data centers. Anecdotal evidence indicates that recently completed data center facilities are being converted to other uses. Other scenarios suggest that computing electrical load may actually decrease as the computational capability of future generations of IT equipment will outstrip the computational need and allow older equipment to be retired. [Anonymous 2001; Baer; Bors 2000; Mandel 2001]. For these reasons, identifying and tracking energy trends in the industry is a prerequisite for coping with increasing energy intensity within data center facilities and to predict the impact on electric utility infrastructure.

Possible Public Interest Actions:

- ◆ Update the California data center market assessment and develop a better understanding of the market by surveying industries that provide specialized goods or services for data centers such as manufacturers of raised floor, or UPS systems.
- ◆ Monitor trends in the data center market, such as space availability and processor heat intensity, through collaboration with industry associations such as iTherm, 7x24 Exchange, and the Uptime Institute
- ◆ Project future data center market and energy demand by working with industry associations
- ◆ Develop market data at the utility level to facilitate system planning and identification of potential bottlenecks. Monitor Utility load requests for new projects

Benchmarking Energy Use

Understanding Data Center Computer Load Density

A comprehensive literature review revealed that there is limited information available for understanding data center energy performance. This situation has created considerable confusion and conflicting information concerning the true electrical power demand in data centers [Baer; Bors 2000; Hellmann 2002; Mandel 2001; Mitchell-Jackson et al. 2001; Roth et al. 2002; The

Uptime Institute 2000; Thompson 2002; Wood 2002]. Appropriate PIER involvement would be to provide an overview of the current energy use through benchmarking a diverse sampling of the state's data centers. This would establish a baseline to develop an understanding of current operation and enable comparison to similar facilities. The benchmarking framework could then be used to track energy performance over time using a consistent set of metrics. As has been demonstrated with other building types and equipment ratings (i.e. EnergyStar), benchmarking will lead to improved energy efficiency through identification and use of best practices in the case of building systems, and improved component efficiency for items such as computer room air conditioners, or computer power supplies. It is also likely that areas requiring research to get over other technological or institutional barriers will be identified.

During our research IT professionals and data center designers expressed a good deal of confusion regarding data center load densities. There is a wide variety of computing and communication equipment each characterized by varying energy demand and intensity. There currently is little measured benchmark data for energy end-use taking into account load diversity and other operational factors. Hence IT professionals and data center designers frequently overstate energy requirements by relying on nameplate ratings or other conservative estimates. IT equipment includes many types of devices from mainframe computers to "blade" servers to disc storage devices yet the problem of identifying the composite true composite electrical load is a common theme.

Even after current load density benchmarks are established, they will likely require continuous maintenance as industry conditions rapidly evolve. The trend in processors and storage media has been to provide exponential improvement in computing capability as predicted by "Moore's Law", and processors have exhibited corresponding increases in heat density. This trend produces locally intense heat at the processor and when servers are stacked together - in the data center. Many data centers are constantly adding and/or removing processing equipment due to growth, changes of occupants, or technology improvements. While these changes have relatively little impact in the short term, they can lead to load growth for the data center over time. This situation leads to difficulty in understanding the operational state for the current collection of IT equipment and the situation becomes even less clear when trying to predict future trends. In one scenario, processor heat load is expected to rise exponentially as it has in the recent past. In another, processors and related components are expected to become more thermally efficient. And in yet another, computing capability is theorized to outstrip computing needs resulting in fewer IT devices. There are also load uncertainty issues due to electrical load diversity such as occurs within computing equipment due to various operational states (sleep mode, full processing, data storage, etc.) and on a macro level for all electrical systems (various operating combinations of IT and infrastructure equipment, fans or compressors on or off, etc.).

By understanding the current heat producing electrical loads, and trending their changes over time, the industry can better design systems to operate efficiently today, and make them adaptable for efficient operation in the future. Limited benchmarking and case studies to date provide insight into the actual range of energy intensity in California data centers (figure 6). However, the load densities exhibited by facilities studied varies widely, and further work remains to characterize these facilities so that the data collected can be used to predict the load

density of future facilities. Additional benchmarking will help to provide comparative data for various types of data centers and are very likely to lead to identification of best practices.




Figure 5 Computer Load Density Benchmark Results

Possible Public Interest Actions:

- ◆ Develop robust benchmark data by compiling available end use data and adding new benchmark data through case studies or industry self- benchmarking.
- ◆ Encourage sub-metering and instrumentation to facilitate monitoring energy end use in data centers.
- ◆ Through case studies, illustrate the margin between actual loads and original design loads.
- ◆ Benchmark actual operating temperature and humidity ranges in data centers.
- ◆ Develop and deploy a benchmarking protocol to enable data center designers, owners and operators, commissioning agents, and other energy engineers to perform benchmarking in a standard manner.
- ◆ Develop a database of energy benchmarks using standard benchmark data collected through case studies or self-benchmarking (by use of standard protocol).

Developing Energy Metrics

Benchmarking, using consistent metrics provides data center industry professionals with valuable data to enable performance comparisons much the way that commercial buildings are compared. Facility designers and operators have traditionally used the metric, kW/sq. ft., as a

basis to define computer, lighting, HVAC, and other electrical loads. However, computer equipment load intensity has been expressed in many different ways based upon:

- ◆ Area under server rack
- ◆ Area of raised floor
- ◆ Total building area
- ◆ Raised floor area, less HVAC equipment occupying space on the raised floor
- ◆ Other variations

Additionally, some data center professionals have abandoned the kW/sq. ft. metric in favor of W/rack, with the number of racks determined from physical space available. Providing a consistent metric to define IT equipment load intensity is important for a consistent understanding of design capacity and actual performance.

Other metrics that quantify computational efficiency such as millions of instructions per second per kilowatt (MIPS/kW) are also being proposed³. Key facility systems' efficiencies can be evaluated through the use of other metrics such as kW/ton of chilled water, or cfm/kW of air, both of which provide direct system level efficiency comparisons.

To develop metrics most useful to the data center market, the first step may involve an examination of the pros and cons of existing metrics used by engineers and researchers from different disciplines [Aebischer et al. 2002a; Aebischer et al. 2002b; Beck 2001; Feng et al. 2002; Mitchell-Jackson 2001; PG&E 2001]. The second step would be to further refine and get consensus on the metrics that can be of most use. One data center industry association, the Uptime Institute (www.uptime.com), has attempted to standardize the definitions of kW/sf in data center spaces, however their target constituency represents only a portion of the data center market, and other important facility system efficiency metrics (such as chilled water system efficiency in kW/ton) are not addressed. Nonetheless their data helps to characterize the current computing equipment load density in data centers (figure 6). On average this data correlates well with case studies performed to date.

³ MIPS is defined as “million instructions per second” and is a measure of the rate that computations are occurring in a computer




Figure 6. Uptime Institute Energy Density Data

Possible Public Interest Actions:

- ◆ Work with industry associations such as 7X24 exchange (www.7x24exchange.org), Uptime Institute (www.uptime.org), Silicon Valley Manufacturers Group (www.svmg.org), AFCOM (www.afcom.com), Advanced Computing Systems Association - USENIX (www.usenix.org) etc. to develop metrics of most use for the data center community.
- ◆ Provide workshops with industry associations and public utilities to promote use of standard metrics
- ◆ Develop protocols that include and integrate energy efficiency into data center performance metrics (e.g., availability, reliability). Develop case studies to evaluate the related construction and operational cost implications of these key areas.

Benchmarking IT Equipment - Actual vs. Nameplate

Predictions of electrical requirements for IT equipment are often determined by use of "nameplate" values. Common nameplate information for most pieces of computer or network equipment usually provide electrical values designed with a "safety factor" to ensure that the equipment will energize and run safely. Typically the values specified by the manufacturer are conservatively set with little correlation to normal operational conditions. When equipment nameplate information is used directly to develop facility power consumption and resulting cooling requirements, the facility systems are often oversized by factors of four or more. Obtaining and publicizing true power demand for IT equipment would provide a much needed, rational basis for determining real power requirements. Comparing actual and nameplate values will provide important insight for IT and facility professionals and can lead to improved sizing of electrical and mechanical systems. One professional described the need to determine an electrical "Expected Maximum Load" (EML) and resulting "Expected Maximum Heat" (EMH) rejected for each piece of equipment as an alternative to nameplate ratings.

Possible Public Interest Actions:

- ◆ Benchmark actual loads of various types of IT equipment typically found in data centers and compare to nameplate values. Publicize findings and develop training guidance to deal with nameplate values in data center facility design.
- ◆ Deploy guidance through workshops and training sessions.
- ◆ Develop testing protocols to characterize the Expected Maximum Load (EML) and Expected Maximum Heat (EMH) by working with manufacturers of IT equipment.
- ◆ Monitor projections from industry organizations such as iTherm (www.itherm.org), Intel's Developer's Forum, etc.
- ◆ Investigate characteristics of new and emerging computer technologies, such as blade servers.

Identifying and Developing Best Practices

By reviewing case studies and benchmark data, the better performing facility systems will become evident. To investigate how better energy performance is achieved, the design features and operating procedures deployed in the top energy performing systems can be examined and documented. In this way, best practices for the data center infrastructure will be identified. For example, superior equipment used in interfacing IT equipment with facility systems (racks, cabinets, etc.) also affects energy performance and similar to the facility systems, better energy efficiency performance in certain configurations can be identified. Examples of best practices in the operating procedures area might include selections of the range of allowable temperature and humidity. . Once best practices are identified, existing and new data centers' performance can be improved by actively promoting the better performing technologies and strategies.



Figure 7 Server Cabinets and Racks

Possible Public Interest Actions:

- ◆ From benchmarking activities, identify the top performing data centers from an energy perspective.
- ◆ Conduct investigations at these facilities to determine which practices contributed to such performance.
- ◆ Confirm the cost-effectiveness of these practices.
- ◆ Disseminate information about the identified best practices to major industry stakeholders.
- ◆ Add to, or modify the energy research roadmap as gaps in available solutions are identified.
- ◆ Determine widest observed temperature and humidity set point operating range in data center spaces. Work with industry associations to establish broader ranges yet maintain reliability.
- ◆ Research available modeling tools and provide designers with comparative data
- ◆ Survey available energy storage/un-interruptible power technologies and their relative efficiencies

Improving Facility Infrastructure Efficiency

Develop Better Tools including Planning and Design Guides

Data Center designers and owners confirm that critical facets of the design and operation of data centers are often accomplished through rough approximations or trial and error. There are many tools currently available for modeling energy, airflow, and heat transfer, however they are not routinely used in data centers. Apparently, the complexity, technical limitations, perceived inaccuracy, and cost of these tools are barring them from widespread use in the data center sector. To improve the accessibility of available tools, research will be required to provide data center designers and operators with a comprehensive base of existing resources. Areas where better design tools are needed could then be identified and targeted for future development. For example, the use of Computational Fluid Dynamics (CFD) software to design HVAC systems for complex arrangements of server racks may be desirable but available software packages are unwieldy in their current state. Research to improve such software may well result in efficiency gains for this industry.




Figure 8 Computational Fluid Dynamics Models

Another example of where better tools are needed is in sizing and placement of floor tiles to provide air-cooling to racks of computers through raised floor systems. Design professionals describe current practice as far from an exact science, and accomplished through judgment, experience, or trial and error. Examining best practices may yield some clues, however the industry needs a simplified, yet accurate method to assist in energy efficient computer room HVAC design.

Over-sizing of electrical infrastructure is epidemic in the data center industry. Accounts of installed infrastructure capable of serving power densities ranging from 100 to 300 Watts per square foot (W/ft^2) have been routinely cited in media reports about these facilities (Stein 2002). Yet, both published and unpublished studies of actual data center power demand suggest that on average most of these facilities actually exhibit energy intensities between 30 to 55 W/ft^2 (Mitchell-Jackson 2001).



Figure 9 Representative Loading in Data Centers

Mitchell-Jackson (2001) and others have identified numerous reasons why data center operators and designers oversize electrical infrastructure. An article in Energy User News illustrated this issue graphically as shown in Figure 9.

The following reasons that designers over-size electrical systems are cited:

- ◆ Actual power requirements for computing equipment are often much less than nameplate data. Nameplate ratings are sometimes used as the basis to size electrical equipment and then result in oversized mechanical HVAC equipment for cooling.
- ◆ They have little guidance regarding how the power demand of electronic equipment varies depending on whether that equipment is in active mode, or is idling. (i.e. diversity of equipment load)
- ◆ They sometimes, inappropriately, apply power densities based on small areas to much broader areas such as power density for a computer rack being applied to the entire raised floor area.

In the last decade, many data centers were hastily planned since IT equipment had to be quickly installed on fast-track schedules to meet business objectives. One Case Study participant related that it was impossible to find people with data center design or operational experience during that time, and that the industry has been learning as they go. Often, undesired consequences induced by the lack of planning and experience were not apparent until after serious reliability problems or poor environmental control had already occurred [Sullivan 2002; Thompson 2002].

Designers are tasked with accurately predicting space, energy requirements, and cooling needs to ensure data center reliability. Due to the fast-track nature of data center projects, the lack of experienced technical expertise, and the myriad of design decisions encountered, it is appropriate to consider a data center planning and design guide addressing efficiency in key areas. Such a tool could include guidance on thermal trends, growth in the amount and intensity of servers, incremental build-out, and flexibility [Anonymous 2001]. Efficiency gains could be maximized by incorporating green building principles and by integrating more efficient IT and facility equipment as they become commercially available. To ensure adoption, energy efficiency, reliability, and security would need equal consideration and evaluation. [Anonymous 2002a; Beck 2001].

Many in the Data Center Industry expect the power density of electronic equipment commonly used in data centers to rise rapidly in the near future (Uptime Institute 2000). As a result, some types of data centers believe that they must demonstrate to their customers that they have surplus capacity and redundant systems - resulting in a "more is better" philosophy - even where additional electrical power is not reasonably going to be needed.

Regardless of the reasons, the over-sizing of electrical infrastructure has several consequences, most of which are negative for both the data center industry and society at large. Some oversized electrical equipment operates inefficiently at small part loads, which wastes energy. Excessive capital costs and possible delays in obtaining power from the local utility are likely outcomes when electrical equipment is oversized and this becomes a barrier to development in the data center sector. Lastly, many electric utilities responding to power requests based upon exaggerated power demand estimates may over invest in transmission and distribution infrastructure, or as has happened in some California locations, may deny the request for service based upon transmission constraints, forcing the data center to be located elsewhere - possibly outside of the state. In addition, some utilities are contemplating rate schedules that include provisions for recapturing capital cost of new transmission and distribution infrastructure.

Possible Public Interest Actions:

- ◆ Research available data center design and analysis tools and summarize their features on a website.
- ◆ Independently confirm adequacy of CFD modeling tools to accurately predict thermal performance.
- ◆ Identify systems, components, and issues for which guidance is needed. Develop a guide (or guides) incorporating current best practices along with any new ideas.
- ◆ Develop advanced design and modeling tools
- ◆ Develop and implement modular (scalable) system concepts to improve part load efficiency
- ◆ Develop mechanical and electrical system sizing guidelines including use of benchmark results to account for load diversity to allow efficient operation initially and as IT equipment loads change. Consider the relationships between reliability, availability, and energy efficiency.
- ◆ Use benchmark data to establish a correlation between the nameplate ratings and actual loads associated with IT equipment. Work with industry associations to influence manufacturers to establish and publish realistic nameplate values for various operational states (i.e. sleep mode etc.)
- ◆ Use benchmark data to develop guidelines to account for the role of equipment diversity (active vs. idle) when estimating data center electrical loads.

Developing Better Monitoring and Control Tools

Although building management systems are currently used to monitor and control energy intensive systems operating cost in the data center sector, they are rarely used to optimize energy performance. Research is needed to develop and deploy improved building monitoring systems which are able to evaluate and correct energy performance as well as protect critical computing equipment.

Possible Public Interest Actions:

- ◆ Survey a sampling of Data Center operators to determine current practice, existing system capability, and barriers to more extensive monitoring and control.

Maintaining Efficiency as IT Equipment Electrical Requirements Change

In order to improve the efficiency of data centers, HVAC and electrical systems are needed that are capable of providing efficient operation over the range of operating loads that occur as data centers are incrementally built-out, or as computing equipment and electrical loads vary with changing technology. Currently, electrical and HVAC systems in most data centers operate considerably below their design basis. Over the life of the facility, many different loading conditions will be present based upon the changing conditions discussed above under benchmarking. The challenge that designers face is to provide design options that satisfy both fully loaded and partially loaded conditions while achieving high efficiency. Design professionals have identified that design options traditionally have been limited by capital budget constraints, unrealistic owner demands, or other reasons, but often are due to lack of good planning guidance. Planning for incremental build-out, for example, could defer the cost of capital equipment until needed in future years, while selecting energy efficient designs matched to smaller initial electrical loads.

Clearly, systems are needed that can maintain efficient operation over extremely wide load variations. Actual electrical loads may vary from design values for many reasons - overly conservative design requirements, change in computing equipment, or simply changes in the mission of the data center. Changes in technology in the future, such as use of smaller, more efficient servers, or use of direct liquid cooling instead of air, may also result in part load operation of conventional cooling systems. Strategies may include incremental connection of UPS systems, chillers, pumps, and fans while using low-pressure drop distribution. Research into optimization strategies and promotion of best practices will increase the likelihood of industry adopting more efficient approaches using current technologies. The same strategies may also enable demand response reductions for emergencies or rate relief. The desire to have more reserve capability may be overcome by demonstrating the ability to economically and quickly increase the capability of infrastructure systems. Future data center business should be more cost-competitive, and designs that can deliver major savings in both capital cost (correct sizing) and operating cost (high efficiency) should provide their owners and operators with a competitive advantage [RMI and DR International 2002].

Possible Public Interest Actions:

- ◆ Develop case studies to demonstrate how modular design of facility systems can improve efficiency and reliability.
- ◆ Develop a model design criteria that facility owners could use to specify efficiency goals.

Electrical System Issues

Power Reliability and Uninterruptible Power Supplies

An uninterruptible power supply (UPS) is used to protect sensitive loads (computers, servers, critical electronic equipment, etc.) from various forms of power disturbances during transmission or distribution that can affect their operation or service life. UPS's are installed between the utility (or distributed) power supply and end-use loads. These systems usually contain both electronic control circuitry and electrical storage. Often, electrical storage is provided by lead-acid batteries, although flywheels seem to be picking up market share.



Figure 10 An older UPS system

Prior benchmarking by LBNL has identified that UPS systems represent a large opportunity for energy savings. UPS systems are continuously energized to provide standby power and power conditioning for IT equipment and varying amounts of facility infrastructure depending upon the critical nature of the center. The UPS energy losses occur due to electrical power conversion and to charge batteries or maintain inertial systems. The efficiency of UPS systems drops off significantly at partial load conditions, which is typically how most data centers operate. See figure 11 below displaying actual measured values

from previous case studies. To determine the loss in the UPS, electrical load was directly measured at the input and output of the UPS.




Figure 11 Measured UPS Performance From Case Studies

To compound the inherent inefficiency of UPS systems, redundancy strategies often call for use of multiple UPS's where each may be lightly loaded.

Even small savings in UPS efficiency can yield large on-going savings. To investigate the efficiency opportunity, research into the available UPS systems' to determine efficiencies over varying percentages of full load is necessary. This research could lead to development of better labeling, more useful ratings, and could be used to develop financial incentives for implementing more efficient systems. Results will provide owners and designers useful information on the efficiency of available systems for various loading conditions. In addition, efficiencies related to common redundancy strategies can be studied to determine best practices for achieving both desired redundancy and energy efficiency. This research has a broader application since UPS systems are also prevalent in other building types such as cleanrooms and laboratories with hazardous environments or where other life safety equipment is involved.

Once best practices utilizing current technology are developed, it should be possible to collaborate with UPS manufacturers and other researchers to develop next generation systems. For example, many researchers are working on improved battery systems, and inertial UPS systems are becoming more prevalent. Collaboration with those developing new battery or inertial technology or may yield breakthroughs in both efficiency and capacity. Other technologies such as fuel cells may also play a role transforming the UPS market.

For facility operators or designers who would like to select UPS based, at least in part, on energy efficiency, the web of options, claims, and counterclaims is confusing at best. Impartial reporting of the efficiency of various systems could provide owners and designers with valuable selection criteria. For example, many major UPS manufacturers claim energy savings associated with their UPS when compared to their competitors. This situation might be improved if there were an independent party that could offer credible advice regarding integrating energy efficiency considerations into the UPS selection process.

Possible Public Interest Actions:

- ◆ Survey available UPS and energy storage technologies. Evaluate controls and control strategies. Determine losses versus load for each UPS
- ◆ Hold workshops involving facility electrical design professionals, and other researchers to investigate UPS system design concepts and configurations to achieve desired redundancy
- ◆ Develop more efficient UPS solutions for various levels of reliability (N+1, N+2, 2N, etc.)
- ◆ Provide training workshops summarizing comparison of various UPS equipment and efficient methods of achieving redundancy.
- ◆ Utilize UPS ratings to develop rebate programs through public utilities
- ◆ Develop a model of data center UPS to evaluate a variety of storage technology combinations. Analyze the lifecycle costs and benefits of the different combinations. Distribute the results of this analysis to data center industry stakeholders.
- ◆ Perform research to develop more efficient energy storage technologies.

Reduce losses due to harmonics

Single-phase electronic loads, such as computers, servers, hard drives, and other devices used in data centers, are susceptible to harmonics, which can cause overheated neutrals in 3-phase/4-wire systems and overheated transformers due to circulating currents. Such overheating can lead to shortened transformer life, as well as potential safety concerns (fires, explosions, etc.) and there are electrical power losses caused by resistance to the large neutral currents. Many types of electronic equipment can generate harmonics. The heating caused by harmonic currents also wastes energy required to provide cooling [Gilleskie 2002; Greenberg 2001; Howe et al. 2001]. Today, commodity servers containing inexpensive power supplies often introduce significant harmonics. When large numbers of servers generating harmonics are present, high transformer operating temperatures result along with high neutral conductor current in three phase electrical systems.

Determining to what extent harmonic currents are induced by clusters computing equipment in data centers could be a worthwhile first step. Then, existing technologies for mitigating these

currents could be evaluated. Based upon these findings, new technologies or strategies to eliminate the source of the problems or to mitigate the effects could be developed.

Possible Public Interest Actions:

- ◆ Research available technologies and develop improved cost effective techniques to diagnose and mitigate harmonics effects
- ◆ Benchmark actual power factor vs. IT equipment manufacturer's claimed power factors

Optimize electrical distribution

There are often several levels of power conversion occurring in data center facilities and within the IT equipment contained therein. Most data centers experience significant electrical power losses in their facility systems' supply and distribution including losses in transformers, power line conditioners, UPS, line losses, etc. Add to these the power losses due to inefficient power conversion (power supplies and/or voltage changes) within IT equipment, and the result is a large loss in useful electrical power coupled with energy use for HVAC systems to remove the heat produced in conditioned spaces.

An optimal system might integrate the IT equipment with the facility in such a way as to minimize power conversions. For example, the individual power supplies in servers could be eliminated if the correct voltages of DC power could be supplied efficiently from a central system, or in the case of fuel cells, directly from the power source. One industry expert envisions the data center of the future similar to a computer in its case. Taking this idea a step further, the electrical system could be thought of as an integrated system from where it enters the data center to the ultimate end use. When viewed in this manner, optimized systems could be designed so as to optimize energy (distribution and conversion losses), reliability, power quality, and potentially provide additional benefits such as elimination of harmful harmonics.

Possible Public Interest Actions:

- ◆ Workshop involving computer/server manufacturers, facility electrical design professionals, and other researchers to investigate system design concepts from the electric meter through to the computing device - whole system approach.
- ◆ Benchmark losses in power conversion throughout the electrical system including transformers losses, UPS losses, power supply losses etc.

Develop integrated solutions to meet demands for power generation, distribution, reliability, and efficiency

Although the focus of this roadmap is on the energy efficiency of the data center and not distributed generation (DG), there are areas where distributed generation may hold promise for efficiency improvements also. Based upon business needs, power reliability and power quality often are the primary driving force in data center design. Mission critical data centers strive to

avoid any sort of power outage because an interruption of power could cost millions of dollars per occurrence.

Grid-connected DG systems and in particular, fuel cell-based systems could be important solutions [Planet-TECH 2002]. One option would be to have the primary power supply be an on-site generation system achieving a minimum of double redundancy, with the grid as backup. For example, use of fuel cells with the grid as backup, could simplify power conversions, be a ready source of uninterruptible power, and could eliminate costly and inefficient equipment.

Strategies can be developed to improve power reliability, power conditioning, uninterruptible power supply (UPS), and improved energy efficiency through use of DG systems [Anonymous 2001]. High availability can also be achieved with Combined Heat and Power (CHP) for Mission Critical Applications [Cratty and Allen 2001].

Possible Public Interest Actions:

- ◆ Demonstrate more efficient power distribution. To simplify the path of power to the computers, the dual redundant, on-line UPS systems could be replaced with simpler self-stabilizing buffer/transient technology systems (flywheels, new high power batteries or super-capacitors), powered by a clean, reliable on-site power source (*e.g.*, turbines, fuel cells, etc.) Part or all of this strategy could be demonstrated in an operating data center.
- ◆ Demonstrate a thermal-based cooling system that uses an on-site generator's waste heat to drive absorption, desiccant or other cooling cycle technology
- ◆ Accelerate the development of reliable, clean, and economically feasible distributed generation technologies (such as fuel cells) for critical power applications

Improve lighting efficiency

Energy used for lighting in data centers represents a small fraction of the overall energy use, yet the opportunity for efficiency savings is great – much more than for commercial office space – due to the amount of time when unoccupied. Modern Internet hosting facilities take advantage of this fact by providing a “lights out” philosophy where lighting is provided only when needed. Standard lighting controls in combination with more sophisticated building management systems can easily achieve a 50% reduction in lighting electrical energy use.

Possible Public Interest Actions:

- ◆ Demonstrate saving potential through cases studies and demonstration projects using existing lighting controls.
- ◆ Develop energy efficient task maintenance lighting to avoid lighting large data center areas for locally small maintenance or installation activity.

HVAC System Issues

Improve the performance of Data Center Air Conditioning

Computer Room Air Conditioners (CRAC units) are the most widely used specialty HVAC components in data centers. The efficiency performance of today's CRAC units operating with raised floor systems, however, is far from optimal. Some examples from case studies of where efficiency was impacted by inefficient CRAC unit deployment and operation illustrate the extent of such problems:

- It was not uncommon to find some CRAC units humidifying while others were de-humidifying the same space.
- Often, CRAC units were not placed in optimal locations.
- In one case, CRAC units were found not to be providing any cooling and could have been turned off, relying on a more efficient central house system.
- Often, more CRAC units were operating than needed.
- Air return to CRAC units did not take advantage of thermal stratification in the data center.
- CRAC units were manually turned on and off.
- Humidification methods were extremely inefficient.
- Openings in raised floors allowed air to bypass its intended use.
- Placement of raised floor tiles with openings was subjective and/or based upon experience with less than optimal results.
- Areas under raised floors were blocked, preventing airflow to where it was needed.

There appear to be many efficiency improvement opportunities with CRAC units and raised floor systems in data centers. These opportunities span design, operations, and maintenance involving the air conditioning units and their controls, the raised floor system, and interaction with other building air systems. Many measures could benefit existing data centers suggesting a public interest role to encourage and provide guidance for energy efficiency commissioning and retrofits. Best practices discovered through case studies and benchmarking, can be applied to both existing data centers and in new data center facilities.



Figure 12 Typical CRAC unit

Possible Public Interest Actions:

- ◆ Develop incentives to apply Best Practices in new or retrofit data center HVAC systems
- ◆ Hold workshops with Industry Associations such as the 7X24 Exchange Organization to develop additional improvements, disseminate best practice information, provide

information on public interest incentive programs provided by the California Energy Commission or Utilities.

- ◆ Develop a demonstration project to illustrate efficiency improvement opportunities, in addition to other benefits such as improved thermal performance.
- ◆ Improve energy efficiency of CRAC units (e.g. more efficient fans, motors, use of variable speed compressors, improved controls, etc.)

Increase the use of “Free Cooling”

As data center HVAC systems typically provide cooling around the clock, every day of the year, they do so during many hours in which the outside air conditions are favorable enough to offset the use of refrigerated cooling. However, anecdotal evidence suggests that few data centers use air-side economizers. As a result, this technology, which has the potential to significantly reduce data center cooling costs, appears to be underutilized.

The reasons for such underutilization are unclear. Perhaps data center designers are concerned that economizers will introduce unwanted humidity or airborne particulates (Shields and Weschler 1998). Perhaps they are concerned that the first and operating costs associated with equipment required to remove humidity and particulates will render such systems uneconomical. This may be further reason to examine the real technological limits of temperature and humidity in data centers. Perhaps they have experienced failure rates similar to those exhibited by economizers in the low-rise commercial buildings industry, and avoid them as a means to maintain high system reliability or lower maintenance costs. These barriers should be explored to determine whether this existing technology could be better utilized. This is one area that may be included for guidance by a newly formed ASHRAE Technical Committee.

Similarly, in many locations, free cooling can be employed to produce chilled water by minimizing the use of chillers. Several different methods can be used to achieve free cooling including direct use of cooling towers, chilled water heat exchangers, and options provided by chiller manufacturers. Although these strategies have been successful in many other chilled water systems, case studies have revealed that free cooling is underutilized in data center applications.

Possible Public Interest Actions:

- ◆ Conduct a study to estimate the costs and benefits associated with using economizers in the data center sector in a variety of climate areas.
- ◆ Survey manufacturers of data center HVAC equipment to determine the market penetration of air-side economizers in data centers.
- ◆ Survey data center operators to identify barriers to more widespread use of air-side economizers.

- ◆ Provide research to develop better economizer technologies and operating techniques which have the potential to increase market penetration.
- ◆ Provide training and raise awareness of free cooling benefits.

Encourage manufacturers of computer room air-conditioner (CRAC) units to release products with variable-speed compressors

Because of uncertainty regarding power densities of electronic equipment, both now and in the future, CRAC units are frequently oversized. As a result, this equipment typically operates to remove extremely low heating loads.

In some applications, a way to improve the efficiency of CRAC units would be to fit them with variable-speed compressors. Such units would have the potential to be four times as efficient at low loads than at rated load. Indeed, CRAC units with variable-speed compressors were developed by a major manufacturer, but never released. The manufacturer determined that few customers would be willing to pay a premium for such a feature. Given that uncertainty regarding equipment power densities has grown since this manufacturer mothballed this product, perhaps some targeted efforts could stimulate a larger market for this device.

Possible Public Interest Actions:

- ◆ Purchase and install CRAC units with variable-speed compressors for one or more showcase projects. Monitor the operation of these units and project their likely costs and benefits. Disseminate the results of these efforts to key stakeholders in the data center industry.

Promote use of highly efficient air delivery systems.

Distribution of air for cooling in a data center can be accomplished in an infinite number of ways, however the paradigm of using raised floor systems for the delivery of air can be problematic in terms of providing cooling where it is needed and with sufficient quantity to cool high heat densities. In addition, energy efficiency can suffer if the lay out of the systems is not optimal, or if unsealed openings are present.

Previous case studies have shown that an HVAC system designed similarly to a traditional building system (similar to those used in large commercial buildings) can be more efficient than the current practice of using raised floor systems with specialized computer room air conditioners.⁴ Sound engineering principles, such as providing large, low-pressure drop delivery systems, along with efficient fans and motors that are controlled for varying load conditions, can be utilized in various configurations. Currently these systems are being successfully used with overhead delivery (and overheard wire management) in a small fraction of data centers. Wider acceptance of this practice, and/or using traditional design practices in concert with raised floor

⁴ Case studies are available through LBNL's data center website: <http://datacenters.lbl.gov>

systems could provide large-scale improvements. Old paradigms where computer room air is recirculated using inefficient air movement and cooling devices need to be challenged. Systems utilizing efficiently sized air handlers and cooling coils, perhaps located outside of the data center could be an attractive alternative to current practice.

Possible Public Interest Actions:

- ◆ Develop efficient system design concepts based upon benchmarking results and good engineering practice. Hold workshops to present the design concepts to design professionals.
- ◆ Demonstrate thermal and energy improvement by optimizing existing raised floor air distribution with an Industry Partner

Determine whether wider ranges of humidity and temperature control can be tolerated by electronics equipment.

Based upon case studies and interviews with designers, it was apparent that many data centers are maintained at approximately 50 percent relative humidity within a tight range, typically plus or minus 5 percent (Nordham, Reiss, and Stein 2001). The need for such tight humidity control is questionable. The reasons given for such control seem to relate to earlier generations of computing equipment where humidified air was needed to prevent static electrical discharges that could damage electronic equipment. In addition, overly moist air could result in water vapor condensing on electronic components, which would also result in equipment damage. These humidity control parameters add energy consumption, both for dehumidifying and/or humidifying. In some cases, dehumidifying is followed by reheat in order to achieve the tight tolerances.

Were humidity levels allowed to fluctuate more, energy consumed to achieve humidity control could be saved. Industry has recognized that this situation is not optimal and is beginning to develop standards and guidelines to enable the design conditions to be relaxed. For example, the manufacturers of electronic equipment used in data centers typically specify relative humidity operating conditions ranging from 20 to 80 percent (Sun Microsystems 2001).

ASHRAE has established a new Technical Committee (TC 9.9) to focus on High Density Electronic Equipment Facility Cooling. This committee is establishing guidelines that should improve the performance of data center facilities. The guidelines focus on several important areas:

- Environmental specifications for various classes of electronic equipment
- Standardized methods of determining temperatures in data centers
- Equipment and facility layout
- Standard reporting of heat load and airflow requirements

These standards, when developed and implemented will help to improve energy efficiency through right-sizing of facility systems, efficient layouts, and broader temperature and humidity limits. Additional research leading to more realistic limits in this area should enable data center designers and operators to control temperature and humidity levels with wider tolerances. Such control will save energy, both for the humidification, and possibly by allowing increased use of outdoor air.

Cooling in a data center provides a workplace environment, and removes heat from critical electronic parts to prolong their life. Looking at these two needs separately may yield some efficiency opportunity. By working closely with IT equipment manufacturers, realistic cooling requirements to protect critical electronic components should be developed. Satisfying true cooling needs may allow relaxation of current practice especially during periods when the data center is not occupied or during periods of peak demand. For example, if even small increases in ambient temperature are acceptable, significant energy savings will result. Often data center operators choose to lower the overall ambient temperature as a solution for “hot spots” in their data center, resulting in overcooling and inefficiency. The ability to tolerate some locally higher temperatures could alleviate this problem.

Possible Public Interest Actions:

- ◆ Participate in the development of data center standards through participation in ASHRAE Technical Committee TC 9.9.
- ◆ Once guidelines or standards are established, provide training for data center designers and operators in California.
- ◆ Use models of data center HVAC systems to estimate the energy savings that could be achieved under a variety of temperature and humidity control scenarios.
- ◆ Survey data center designers and equipment manufacturers to research their basis for current data center environmental specifications.
- ◆ Contact electronics manufacturers to determine the basis for current temperature and humidity recommendations.

Making Better Use of Existing Energy Efficiency Guidance



Figure 13 Centrifugal Chiller

Selecting more efficient equipment is but one step in optimizing facility system performance and should be included along with more comprehensive system measures. Much of the energy intensive infrastructure equipment in data centers such as chillers, pumps, motors, and transformers are common in other building types. Much information concerning efficiencies of this equipment exists yet case studies and other anecdotal information highlight that data center designers and building owners need to be more exposed to information for designing systems and specifying efficient equipment. Existing guidance for efficient system design, such as that provided

by ASHRAE, Cooltools, or DOE's Motor Challenge as well as comparative manufacturer's performance data is under-utilized in the specialized field of data center design.

Possible Public Interest Actions:

- ◆ Develop and provide training workshops for design and operations professionals at several locations throughout California concerning use of existing energy efficiency design information.
- ◆ Publish a list of available related energy efficiency resources on public websites.
- ◆ Develop the basis for incentives for use by Public Utilities to stimulate use of more efficient facility systems and equipment, and specialty components used in data centers.

Improving the Interface Between Building Systems and IT Equipment

Manage computing load and demand

Electrical power for computing equipment and resulting heat loads are relatively constant in data centers as long as the computing equipment set does not change. Unfortunately, computer room equipment is frequently changing due to changes in technology, computing and storage capability, and other business factors. This changes the overall electrical load (and resulting heat load) for the data center and may create the need for redistributing cooling air to relieve local "hot spots". Research into better low cost systems for monitoring and tracking changing electrical loads as well as strategies for allowing systems to respond efficiently to increases or decreases in overall or local loading is needed. One innovative data center professional suggests

shifting computing to other less energy intensive areas of the data center - or - even shifting computing to other geographic locations. The ability to sense and respond to local area "hot spots" will enable this concept to succeed.

Possible Public Interest Actions:

- ◆ By working with leading industry partners, determine strategies to optimize infrastructure systems, sensing and controls, operating systems, and IT Equipment to respond efficiently to computing demands.
- ◆ Develop a demonstration project in an operating data center to prove the feasibility of concepts to sense and redirect computing when local area hot spots are detected.

Develop More Thermally Efficient Racks and Enclosures

Efficient air movement inside data centers is critical to energy performance of these facilities. Many racks used to stack IT equipment in data centers are designed to maximize the number of devices that can be accommodated. Many have design features to make them visually attractive or provide security such as locked glass panel fronts. The result is often a closed, densely packed rack system that does not allow for efficient airflow or worse, can contribute to heat transfer problems with the equipment contained in the racks. Racks and equipment enclosures that better facilitate the flow of cooling fluids across hot components would likely improve the efficiency of data center cooling systems. For example, more efficient cooling using air, chilled water, or other mediums may be possible by integrating the design of rack systems with the equipment they contain, as well as with HVAC and electrical infrastructure. The fragmented nature of the market however makes gains in this area difficult. Computer and other IT equipment manufacturers deal strictly with heat transfer issues within their equipment's envelope. Rack manufacturers develop their products independently from the IT equipment suppliers, and the facility systems. An appropriate role for public goods research would be to investigate methods to facilitate the integration of these disparate elements.

Possible Public Interest Actions:

- ◆ Convene workshops bringing the three different areas (computing equipment, rack, building systems) together with a goal of developing an optimized, integrated solution.
- ◆ Develop incentive programs to reward integrated solutions that are more efficient.

Determine Whether Increased Heat Density Makes Economic Sense

Although a guiding paradigm in the IT industry is that customers want greater functional density, the time has come to question whether further increases in density are worth their cost in energy and cooling infrastructure. Open configurations and spatial distribution of IT equipment heat loads (utilizing more real estate) may be better options than development of increasingly complex air cooling systems. There will be break even points where further increases in density are not cost effective compared with spreading out the heat sources. Evaluating options will involve present and future real estate and energy costs, the cost of HVAC systems, and other

factors but should also consider the relative energy efficiency of cooling systems for a given solution.

Possible Public Interest Actions:

- ◆ Develop “heat intensity” economic evaluation tools considering expected heat intensity, energy use and electrical power rates, facility and infrastructure cost, etc.

Limits of Cooling mediums

The different cooling mediums (air, water, dielectric fluids) which could be employed in a data center, each have limits for effectively removing heat from computing equipment. Determining limits of the various cooling technologies that could be applied to electronic computing equipment and their scalability would be useful. Some research in this area by CEETHERM and other industry associations, such as iTherm are pursuing this research. The impact on building systems for various cooling mediums is not well understood. An appropriate role for public interest research would be to investigate the efficiency implications to the facility for the various mediums and encourage the use of efficient technologies. This would likely involve a collaboration between industry researchers and developers of new innovative products. Applying the results of studies focused on cooling limits at the processor or computer level, to the impact on building systems would provide useful information in planning and optimizing building systems.



Figure 14 Typical Computer Racks

Possible Public Interest Actions:

- ◆ Collaborate with CEETHERM (<http://www.me.gatech.edu/me/publicat/brochures/Mettl/Bro0302.htm>, iTherm (www.itherm.org), and other industry associations to assess the limits of current and emerging technology.

Evaluate Emerging IT Rack Systems and Specialty Products

Georgia Tech University and other research organizations are studying heat transfer in IT equipment and computer rack systems, and limited studies are underway. Emerging new products such as new computer rack systems that incorporate plenums and fans to direct cool air to the computing equipment, are being researched or developed by others. California can benefit by collaborating with these organizations to anticipate important trends in IT equipment’s increased heat intensity and the related impact on electrical power demands. Increasing heat

density will dictate the need to develop more sophisticated cooling solutions. Typically, energy efficiency is not a focus in the development of new products and solutions so long as the devices are cooled adequately. Therefore, an appropriate role for public interest research would be to ensure that energy efficiency is considered while new thermal solutions are being developed. .

Possible Public Interest Actions:

- ◆ Collaborate with CEETHERM (Georgia Tech University) and other researchers. Participate in the iTherm conference – an industry association focused on heat in electronics.
- ◆ Evaluate Emerging Systems and Develop Case Studies to Demonstrate Energy Efficiency

Development of Direct Cooling Systems

Based on prior benchmarking studies, it's evident that HVAC systems can be responsible for half (or more) of the energy consumed by data centers. Much of that energy is used to power fans, which force air through underfloor plenums, and draw air through electronic equipment. Roughly half the energy consumed by data center HVAC systems goes towards powering those fans and cooling the waste heat associated with their operation (Westphalen and Koszalinski 1999). Current HVAC systems exhibit five shortcomings that threaten their continuing dominance in the data center sector:

- The on-board fans incorporated into electronic equipment are especially inefficient and take up room that could be used to house electronic components.
- Air is not an efficient medium for heat transfer. Liquids can move heat much more efficiently.
- The same medium currently used to cool equipment (air cooling) is also used to cool workers. As a result, overcooled workers are endemic in this industry.
- As cool air passes through the IT equipment and rises from the raised floor it is warmed. One study found that air 6 feet off a raised floor was nearly 20°F warmer than air in the plenum. Warm air diminishes the lifetime of the equipment located near the top of racks and cabinets (Schmidt 2001a). The Uptime Institute has found that equipment located at the top of racks is more prone to failure and exhibits shorter life.
- Raised floor forced-air systems are incapable of cooling electronic equipment with a power density that exceeds 150 W/ft² (Schmidt 2001b).

These shortcomings can be largely mitigated by changing to systems that apply cooling directly to hot electronic components. These systems are typically based on one of two approaches. They either spray refrigerant directly on hot electronic components (Shaw et al 2002, Isothermal Systems Research 2002) or they feature other innovative heat removal features such as micro-refrigeration systems installed directly on those components (Anderson). Using these techniques, a much smaller and simpler forced-air system can then be used to condition the space around the

electronic components. While the main driver for such devices is to protect the electronic components, there are also opportunities for energy efficiency gains. An appropriate public interest activity would be to promote energy efficient solutions. The costs and benefits of applying such systems in data centers in a widespread manner are not yet well understood, however.

Possible Public Interest Actions:

- ◆ Survey the state of development of direct cooling technologies for electronic equipment and the relative energy efficiency of each.
- ◆ Investigate the costs and benefits of these technologies.
- ◆ Identify and address barriers to the widespread deployment of these technologies.

Improving Efficiency of IT Equipment

Benchmark Servers for Energy Performance

Not long ago, manufacturers began to release computer servers with low-power microprocessors. These new products allow data centers to pack servers far more densely into a given volume, without appreciably pushing up energy consumption or cooling loads (RLX Technologies 2001). Although, they are now produced by several of the dominant computer manufacturers, such as IBM, Compaq, and Dell, as well as a few start-up manufacturers such as RLX, these low-power servers still constitute a small, but growing, portion of the overall server market (Gartner Dataquest 2002). A typical blade server featuring a low power microprocessor is shown below in figure 14.



Figure 15 Typical Blade server
Source: RLX Technologies

One barrier to the widespread diffusion of low-power servers into the data center sector is that information technology specifiers have no reliable way to compare the *in situ* performance of competing servers. There are widely accepted standards for benchmarking servers for processing speed, but it's not clear how well these standards predict actual performance in a variety of data center environments. These standards only measure processing performance, and there is no standard method to measure input power. The tests on low-power servers have been done only by the manufacturers themselves, which

has led to much squabbling in the industry over the methods they used and the reliability of their results. This is one area that the newly formed ASHRAE Technical Committee is working to address to standardize power measurement and reporting. Lastly, because the processors used in low-power servers operate at lower speeds than top-of-the-line servers, there is no consensus on

how many low-power servers are required to perform a given task that could be performed by fewer, faster, high-power servers (Compaq 2002). The net result is that specifiers are confused, and few take the trouble to investigate the opportunities to reduce operating costs.

Possible Public Interest Actions:

- ◆ Survey benchmarks used for comparing servers to determine their usefulness for comparing energy performance.
- ◆ Survey IT specifiers to determine what information, and in what format, would empower them to compare competing servers on the basis of energy performance.
- ◆ Identify stakeholders in the server benchmarking field and obtain their cooperation in an effort to develop an energy performance benchmark.
- ◆ Develop a new standard benchmarking method for comparing the energy performance of competing servers.
- ◆ With the cooperation of industry stakeholders, encourage the development of an independent entity that could test servers to the new energy performance guidelines.

Improve Reliability and Efficiency of Power Supplies

Power supplies convert AC electrical power to DC at various voltages. In data center equipment such as servers they are usually located inside of the electronic product (internal). Power supply efficiency levels of 80 to 90% at full load are readily achievable in most internal and external power supplies at modest incremental cost through improved integrated circuits and better designs [Calwell and Reeder 2002]. With wide variations in efficiency for similar products, careful selection of energy efficient power supplies is important. Comparative information, although not currently available, would be helpful for IT professionals and facility designers to make informed decisions.




Figure 16 Computer Power Supply Efficiencies

There are often several levels of power conversion occurring in data centers and telecom facilities. Currently, most data centers have significant electrical power losses in their facility systems' supply and distribution including losses in transformers, power line conditioners, UPS, etc. Add to these, the power losses due to inefficient power supplies at the computer, and a significant efficiency opportunity results if an optimal central system can be developed. In addition, current practice introduces significant harmonics in many cases.

Alternatives to current practices to

deliver appropriate power to the computing device should be investigated by looking at the entire electrical supply chain. This could involve conversion of the main power to DC, elimination of individual small power supply devices, etc.

Possible Public Interest Actions:

- ◆ Benchmark and assess currently available technology for power supplies commonly used in IT equipment. Identify best performers.
- ◆ Benchmark losses in power conversion throughout the entire electrical system including transformers losses, UPS losses, power supply losses etc. with a goal of identifying improvements in efficiency for electrical power delivered to its end use.
- ◆ Hold workshops involving computer/server manufacturers, facility electrical design professionals, and other researchers to investigate system design concepts from the electric meter through to the computing device - whole system approach.
- ◆ Develop standards for reporting power supply performance. Through participation in the ASHRAE committee on High Density Electronics Facility Cooling it may be possible to introduce such a standard. Alternatively, develop a non-proprietary rating system administered by an independent organization.
- ◆ Encourage development of more efficient power supplies:
 - To reduce No load loss.
 - To improve standby performance.
 - Use two different size power supplies to maximize efficiency for a given load

IT Hardware and Software

A number of strategies dealing with IT equipment and software were identified during the roadmap development. Many of the research areas are promising and could yield dramatic energy performance improvement. The specialized and highly technical nature of most of the ideas however, did not seem appropriate for public interest research at the state level. Much research is on going by various industry players in anticipation of higher heat density, yet the motivating force usually is not energy efficiency. It is important to be aware of developments in related research areas to reward or encourage energy saving innovations. Even though development will be performed by the industry itself for issues such as better sleep modes, or passive heat transfer devices, public interest programs can provide a vital role in demonstrations and market transformation activities. Encouraging low-power computing is a complex topic. Makers of highly efficient computer chips, for example, claim their product can cut energy use significantly (e.g., in half), while others point out a limited application of the chips [Anonymous 2001].

Possible Public Interest Actions:

- ◆ Foster new and emerging technologies by demonstrating energy saving devices, systems, or strategies. Possible candidates for demonstration projects include:
 - o Direct Spray Cooling of Refrigerant onto computer chips
 - o Efficient heat sinks
 - o Direct Refrigeration of processors
 - o Software to redirect computing to eliminate thermal spikes on chips
 - o Demonstration of monitoring and control to match cooling to varying heat load in localized areas
 - o Demonstration of energy efficient computer cabinets
 - o Demonstration of alternate cooling media (N2 cryogenic, spray cooling,) single and 2-phase systems, one fluid from building to chip
 - o Demonstration of best practice for idle mode performance

Appendix A

Organizations Associated with Data Centers

Many organizations are stakeholders in the efficient operation of data centers. They include industry associations, individual industrial companies, research organizations, public interest groups, utilities, DOE, EPA, etc. Web sites for selected organizations are listed below:

iTherm

iTherm is an association of computer manufacturers, semiconductor manufacturers, and researchers focusing on heat removal issues at the chip and computer level.

www.itherm.org

CEETHERM

CEETHERM is a collaborative research program carried out at the University of Maryland and Georgia Tech University. Its focus is to research efficiency improvements for data centers from the chip level through to the building systems including distributed generation applications.

<http://www.me.gatech.edu/me/publicat/brochures/Mettl/Bro0302.htm>

7X24 Exchange

The 7X24 Exchange is an industry association who's goal is to improve end to end reliability by promoting an exchange of information between those that design, build, and maintain data center facilities.

www.7X24exchange.org

Uptime Institute

The Uptime Institute focuses on improving uptime management in Data Center Facilities and Information Technology organizations. Members represent Fortune 500 companies who collectively and interactively learn from each other. They sponsor meetings, tours, benchmarking, best practices, uptime metrics, and provide abnormal incident collection and analysis. They also provide seminars, training in IT, and facilities site infrastructure uptime management, and conduct sponsored research.

www.uptime.com

AFCOM

AFCOM is an association for data center professionals offering services to help support the management of data centers around the world. AFCOM was established in 1981 to offer data center managers the latest information and technology through annual conferences, published magazines, research and hotline services, and industry alliances.

www.afcom.org

Electric Power Research Institute

The Electric Power Research Institute has investigated various aspects of data centers including use of distributed power, power quality, and reliability. They have participated in the Consortium for Electric Infrastructure to Support a Digital Society (CEIDS)

www.epri.com

Silicon Valley Manufacturers Group

The Silicon Valley Manufacturers Group is an industry association in the Silicon Valley area of California. Many data center operators are members of this organization, which generally focuses on improving business competitiveness of the member companies. This focus includes interest in improving reliability and efficiency of operations.

www.svmg.org

United States Telecom Association

The United States Telecom Association is dedicated to serving the needs of the telecommunications community. They provide a forum for the telecommunications industry to define issues, debate policy and set goals. They provide lobbying, and regulatory interface for the industry.

www.usta.org

References

ACEEE, and CECS. 2001. Funding prospectus for "Analysis of Data Centers and their implications for energy demand". Washington, DC, American Council for an Energy Efficient Economy (ACEEE); Center for Energy and Climate Solutions (CECS). July 2001.

The paper includes an overview of data centers; discusses energy use, energy choices, and energy efficiency in data centers; potential impacts of data centers; present and future regulatory issues; and business opportunities in energy services.

Aebischer, B., R. Frischknecht, C. Genoud, A. Huser, and F. Varone. 2002a. Energy- and Eco-Efficiency of Data Centres. A study commissioned by Département de l'intérieur, de l'agriculture et de l'environnement (DIAE) and Service cantonal de l'énergie (ScanE) of the Canton of Geneva, Geneva, November 15.

The study investigates strategies and technical approaches to fostering more energy-efficient and environmentally sound planning, building and operating of data centres. It also formulate recommendations on how to integrate the findings in the legal and regulatory framework in order to handle construction permits for large energy consumers and promote energy efficiency in the economic sectors. Seventeen recommendations grouped in four topics are derived from study conclusions: Transfer of the accord into an institutionalized legal and regulatory framework; Energy-efficiency policies for all large energy consumers; Preconditions, and prerequisites; Operational design of voluntary energy policies.

Aebischer, B., R. Frischknecht, C. Genoud, and F. Varone. 2002b. Energy Efficiency Indicator for High Electric-Load Buildings. The Case of Data Centres. Proceedings of the IEECB 2002. 2nd International Conference on Improving Electricity Efficiency in Commercial Buildings. Nice, France.

Energy per unit of floor area is not an adequate indicator for energy efficiency in high electric-load buildings. For data centres we propose to use a two-stage coefficient of energy efficiency $CEE = C1 * c2$, where $C1$ is a measure of the efficiency of the central infrastructure and $c2$ a measure of the energy efficiency of the equipment.

Anonymous. 2001. Model Data Center Energy Design Meeting. Austin Energy, Austin, TX, Feb 12-13. http://www.austinenergy.com/business/energy_design_meeting.htm

Anonymous. 2002a. 7 x 24 Update: Design & Construction - Issues and trends in mission critical infrastructure design, planning and maintenance.

<http://www.facilitiesnet.com/BOM/Jan02/jan02construction.shtml>. July 23, 2002.
<http://www.7x24exchange.org/>.

Anonymous. 2002b. Continuous Availability Review (CAR). The Uptime Institute: Computersite Engineering, Inc. <http://www.upsite.com/csepages/cseccar.html>. July 22, 2002.

Anonymous. 2002c. End-to-End Reliability Begins with the User's Definition of Success. The Uptime Institute. <http://www.upsite.com/TUIpages/editorials/endtoend.html>. July 22, 2002.

Anonymous. 2002d. Mechanical Systems Diagnostic Review (MSDR). The Uptime Institute: Computersite Engineering, Inc. <http://www.upsite.com/csepages/csemsdr.html>. July 22, 2002.

- Anonymous. 2002e. Site Infrastructure Operations Review (SIOR). The Uptime Institute: Computersite Engineering, Inc. <http://www.upsite.com/csepages/cseior.htm> l. July 22, 2002.
- Baer, D. B. Emerging Cooling Requirements & Systems in Telecommunications Spaces, Liebert Corporation.
During the last several years, power density trends, and consequently thermal density trends in telecommunications spaces have become topics of increasing interest. This paper identifies several of the underlying drivers of these trends, project possible outcomes, and assess the impact on cooling system design for these spaces.
- Beck, F. 2001. Energy Smart Data Centers: Applying Energy Efficient Design And Technology To The Digital Information Sector. *Renewable Energy Policy Project (REPP)*: Washington, DC. (November 2001 REPP).
http://www.repp.org/articles/static/1/1036512059_982708661.html
Both utilities and data center owners face challenges in meeting electricity demand loads with required levels of reliability. However, the bursting of the high-tech stock bubble in 2000 and the 2001 U.S. economic downturn has slowed expansion of data centers. This provides time and an opportunity to examine data center construction and operational practices with an eye toward reducing their energy demands through use of energy efficient technologies and energy smart design practices. As the economy recovers and the next data center rush approaches, best practices can reduce energy use while maintaining or even increasing data center reliability. Energy demands of data centers that support the digital information- and communications-based economy need not be as high as some predict. In fact, data center power demands could be reduced by 20 percent with minimal efficiency efforts, and by 50 percent with more aggressive efficiency measures.
- Blount, H. E., H. Naah, and E. S. Johnson. 2001. Data Center and Carrier Hotel Real Estate: Refuting the Overcapacity Myth. Lehman Brothers: TELECOMMUNICATIONS, New York, June 7, 2001. <http://www.lehman.com>
An exclusive study examining supply and demand trends for data center and carrier hotel real estate in North America. Lehman Brothers and Cushman & Wakefield have completed the first in a regular series of proprietary studies on telecommunications real estate (TRE), including carrier hotels and data centers.
- Bors, D. 2000. Data centers pose serious threat to energy supply. *Puget Sound Business Journal (Seattle)* - October 9, 2000.
<http://seattle.bizjournals.com/seattle/stories/2000/10/09/focus5.html>
To cope with increasing energy demand from data centers, the author discussed feasibilities of two possible approaches: 1) energy industry approach by looking at alternative energy supply 2) construction industry approach by looking at data center energy efficiency. To get there, it is worth investigating five distinct components: (I) Co-generation of power. Presently, standby diesel generators are required to maintain the desired level of reliability at most data center sites, but their exhaust makes most of these generators unacceptable for long-term power generation; (II) Fuel cells offer the promise of very clean emissions and the reasonable possibility for use as standby power; (III) Increased efficiency in data center power distribution systems. There are two separate items that are major contributors to data center power distribution system inefficiencies. The first, power distribution units (PDUs) are available with optional internal transformers that use less energy than the present cadre of K-rated transformers. The second, uninterruptible power systems (UPSs) come in a range of efficiency ratings. If the use of high-efficiency PDUs and UPSs are combined, they offer the potential of a 6 percent saving (IV) Increased efficiency in mechanical cooling systems. In order to ensure data center reliability, mechanical equipment is often selected as a large number of small, self-contained units, which offers opportunities to improve efficiencies; (V) Reductions in energy use by computer, network and storage equipment. Computer manufacturers can do their part by creating computers with greater computational power per watt. They

have been doing this for years as a side effect of hardware improvements, and they can do even better if they make it a goal.

Brown, E., R. N. Elliott, and A. Shipley. 2001. Overview of Data Centers and Their Implications for Energy Demand. Washington, DC, American Council for an Energy Efficient Economy, Center for Energy & climate Solutions (CECS). September 2001.

<http://www.aceee.org/pdfs/datacenter.pdf.pdf>

The white paper discusses data center industry boom and energy efficiency opportunities and incentives in Internet data centers. Emerging in the late 1990's, data centers are locations of concentrated Internet traffic requiring a high-degree of power reliability and a large amount of power relative to their square footage. Typically, power needs range from 10-40MW per building, and buildings are typically built in clusters around nodes in the Internet fiber-optic backbone. During the development boom in 1999 and 2000, projects averaged 6-9 months from site acquisition to operation, and planned operational life was 36 months to refit. Even high energy-prices were dwarfed by net daily profits of 1-2 million dollars per day for these buildings during the boom, creating little incentive for efficient use of energy.

Callsen, T. P. 2000. The Art of Estimating Loads. *Data Center* (Issue 2000.04).

This article discusses the typical Data Center layout. It includes floor plan analysis, HVAC requirements, and the electrical characteristics of the computer hardware typically found in a Data Center.

Calwell, C., and T. Reeder. 2002. Power Supplies: A Hidden Opportunity for Energy Savings (An NRDC Report). Natural Resources Defense Council, San Francisco, CA, May 22, 2002. <http://www.nrdc.org>

The article discusses the efficiency of power supplies, which perform current conversion and are located inside of the electronic product (internal) or outside of the product (external). The study finds that most external models, often referred to as "wall-packs" or "bricks," use a very energy inefficient design called the linear power supply, with measured energy efficiencies ranging from 20 to 75%; that most internal power supply models use somewhat more efficient designs called switching or switch-mode power supplies; and that internal power supplies have energy efficiencies ranging from 50 to 90%, with wide variations in power use among similar products. Most homes have 5 to 10 devices that use external power supplies, such as cordless phones and answering machines. Internal power supplies are more prevalent in devices that have greater power requirements, typically more than 15 watts. Such devices include computers, televisions, office copiers, and stereo components. The paper points out that power supply efficiency levels of 80 to 90% are readily achievable in most internal and external power supplies at modest incremental cost through improved integrated circuits and better designs.

Compaq. 2002. *Compaq ProLiant BL 10e Delivers Industry Defining Transactions per Watt and Transactions per Square Foot*. White Paper.

ftp://ftp.compaq.com/pub/products/servers/benchmarks/BL10e_webbench.pdf

Cratty, W., and W. Allen. 2001. Very High Availability (99.9999%) Combined Heat and Power for Mission Critical Applications. *Cinintel 2001*: 12. <http://www.surepowersystem.com>

Elliot, N. 2001. Overview of Data Centers and their implications for energy demand. Washington, DC, American Council for an Energy Efficient Economy. Jan 2001, revised June 10, 2001.

Feng, W., M. Warren, and E. Weigle. 2002. The Bladed Beowulf: A Cost-Effective Alternative to Traditional Beowulfs. *Cluster2002 Program*. <http://www-unix.mcs.anl.gov/cluster2002/schedule.html>; <http://public.lanl.gov/feng/Bladed-Beowulf.pdf>

Authors present a novel twist to the Beowulf cluster - the Bladed Beowulf. In contrast to traditional Beowulfs, which typically use Intel or AMD processors, the Bladed Beowulf uses Transmeta processors in order to keep thermal power dissipation low and reliability and density high while still achieving comparable performance to Intel- and AMD-based clusters. Given the ever-increasing complexity of traditional super-computers and Beowulf clusters; the issues of size, reliability, power consumption, and ease of administration and use will be "the" issues of this decade for high-performance computing. Bigger and faster machines are simply not good enough anymore. To illustrate, Authors present the results of performance benchmarks on the Bladed Beowulf and introduce two performance metrics that contribute to the total cost of ownership (TCO) of a computing system - performance/power and performance/space.

Frith, C. 2002. Internet Data Centers and the Infrastructure Require Environmental Design, Controls, and Monitoring. *Journal of the IEST* **45**(2002 Annual Edition): 45-52.
Internet Data Centers and the Infrastructure Require Environmental Design, Controls, and Monitoring. The author points out that specifications and standards need to be developed to achieve high performance for mission-critical Internet applications.

Gartner Dataquest. 2002. "Gartner Dataquest Chops Industry's Rapid Growth Expectations for Blade Servers." Press Release.
http://www4.gartner.com/5_about/press_releases/2002_02/pr20020205d.jsp

Gilleskie, R. J. 2002. The Impact of Power Quality in the Telecommunications Industry. Palm Springs, CA, June 4. <http://www.energy2002.ee.doe.gov/Facilities.htm>
The workshop addresses the unique issues and special considerations necessary for improving the energy efficiency and reliability of high-tech data centers. This presentation addresses impacts of power quality including voltage sags, harmonics, and high-frequency grounding in telecommunication industry.

Grahame, T., and D. Kathan. 2001. Internet Fuels Shocking Load Requests. *Electrical World* **Vol. 215** (3): 25-27. http://www.platts.com/engineering/ew_back_issues.shtml
This article discusses the implications of the increase for power demand by the Internet's traffic growth on utility planning, operation, and financing.

Greenberg, D. 2001. Addendum to ER-01-15: A Primer on Harmonics. E-SOURCE, Boulder, Colorado, September 2001.
The electrical distribution systems of most commercial and industrial facilities were not designed to operate with an abundance of harmonics-producing loads. In fact, it is only within recent years that such loads have become widespread enough for industry to take notice and to begin to develop strategies to address the problems that harmonics can create. By 1992, concern about the issue had grown sufficiently that the Institute for Electrical and Electronic Engineers (IEEE) developed and published its standard 519, "IEEE Recommended Practices and Requirements for Harmonic Control in Electrical Power Systems," which established an approach for setting limits on the harmonic voltage distortion on the utility power system and on the harmonic currents created individual power consumers. Since that time, the electronic loads that give rise to harmonic currents have grown dramatically and are projected to continue growing for the foreseeable future. This being the case, there is and will continue to be a market for technological solutions to the problems that harmonics can cause.

Gross, P. 2002. Needed: New Metrics. *Energy User News*.
http://www.energyusernews.com/eun/cda/articleinformation/features/bnp_features_item/0,2584,82741,00.html

Gruener, J. 2000. Building High-Performance Data Centers. *Dell Magazines - Dell Power Solutions* (Issue 3 "Building Your Internet Data Center").

http://www.dell.com/us/en/esg/topics/power_ps3q00_1_power.htm;
http://www.dell.com/us/en/esg/topics/power_ps3q00-giganet.htm

The introduction of Microsoft SQL Server 2000 is a milestone in the race to build the next generation of Internet data centers. These new data centers are made up of tiers of servers, now commonly referred to as server farms, which generally are divided into client services servers (Web servers), application/business logic servers, and data servers supporting multiple instances of databases such as SQL Server 2000.

Hellmann, M. 2002. Consultants Face Difficult New Questions in Evolving Data Center Design. *Energy User News*.

http://www.energyusernews.com/CDA/ArticleInformation/features/BNP_Features_Item/0,2584,70610,00.html

While few data center design projects are alike, there are always the twin challenges of "power and fiber." And sometimes, even local politics and human factors. The paper suggested that the consultant should be brought in as soon as a business case is established so criteria can be established and a concept can be developed, priced, and compared to the business case. A planning is necessary before moving on to site selection and refine the concept and again test the business case.

Howe, B., A. Mansoor, and A. Maitra. 2001. Power Quality Guidelines for Energy Efficient Device Application - Guidebook for California Energy Commission (CEC). Final Report to B. Banerjee, California Energy Commission (CEC).

Energy efficiency and conservation are crucial for a balanced energy policy for the Nation in general and the State of California. Widespread adaptation of energy efficient technologies such as energy efficient motors, adjustable speed drives, improved lighting technologies will be the key in achieving self-sufficiency and a balanced energy policy that takes into account both supply side and demand side measures. In order to achieve the full benefit of energy efficient technologies, these must be applied intelligently, and with clear recognition of the impacts some of these technologies may have on power quality and reliability. Any impediment to the application of these energy efficient technologies by the customers is not desirable for the overall benefit to energy users in California. With that in mind EPRI and CEC has worked to develop this guidebook to promote customer adaptation of energy efficient technologies by focusing on three distinct objectives. 1) Minimize any undesirable power quality impacts of energy-saving technologies; 2) Understand the energy savings potential of power quality-related technologies. These include: Surge Protective Devices (SPDs) or Transient Voltage Surge Suppressors (TVSS), Harmonic Filters, Power Factor Correction Capacitors, Electronic Soft Starters for Motors; and 3) How to evaluate "black box" technologies

Intel. 2002. Planning and Building a Data Center - Meeting the e-Business Challenge. Intel Corp. http://www.intel.com/network/idc/doc_library/white_papers/data_center/. Aug 01, 2002.

The paper discusses the keys to success of Internet Service Providers (ISPs) that include 1) Achieve the economies of scale necessary to support a low price business model; 2) Offer added value, typically in the form of specialized services such as applications hosting to justify a premium price. This document provides a high-level overview of the requirements for successfully establishing and operating an Internet data center in today's marketplace. It offers some of the key steps that need to be taken, including project definition, prerequisites and planning. In order to construct a data center that can meet the challenges of the new market, there are three basic areas of data center definition and development: 1) Facilities: including building, security, power, air-conditioning and room for growth; 2) Internet connectivity: performance, availability and scalability; 3) Value-added services and the resources to support their delivery: service levels, technical skills and business processes. The aim is to provide customers with the physical environment, server hardware, network connectivity and technical skills necessary to keep Internet business up and running 24 hours a day, seven days a week. The ability to scale is essential, allowing businesses to upgrade easily by adding bandwidth or server capacity on demand.

- Koplin, E. 2000. Finding Holes In The Data Center Envelope. *Engineered Systems* (September 2000).
http://www.esmagazine.com/CDA/ArticleInformation/features/BNP_Features_Item/0,2503,8720,00.html
 The paper addresses importance of environmental control in data center facilities. Maintaining data center availability requires absolutely reliable infrastructure. A significant amount of this is devoted solely to maintaining stable environmental parameters. And only constant, thorough regulation and testing of these parameters ensures the integrity of the data center “envelope.”
- Mandel, S. 2001. Rooms that consume - Internet hotels and other data centers inhale electricity. *Electric Perspectives* Vol. 26 (No.3).
http://www.eei.org/ep/editorial/Apr_01/0401ROOM.htm
 The article estimated that the amount of this data center space in the United States nearly doubled in 2000, totaling between 19 million and 25 million square feet by year-end, according to investment analysts. They say they expect another 10 million to 20 million square feet of new space to be added in 2001. Developers are asking electric utilities to supply the buildings with 100-200 watts of electricity per square foot. Since these data centers are new to the economy, there is little historical data on which to base estimates of electricity use for a facility. In addition, the dot.com world makes it difficult for the developer to say confidently how much electricity one of these Internet hotels will use. Source One estimates that tens of billions of dollars worth of electric infrastructure improvements will be needed for data centers over the next few years and that they will consume billions of dollars more worth of electricity. The energy costs are as high or higher than the actual lease costs. Indeed, 50-60 percent of the cost of building a data center is for the power, including batteries, backup generators, and air-conditioning, as well as the cost for utility construction.
- Mitchell-Jackson, J. 2001. Energy Needs in an Internet Economy: A Closer Look at Data Centers, July 2001. <http://enduse.lbl.gov/projects/infotech.html>
 This study explains why most estimates of power used by data centers are significantly too high, and gives measured power use data for five such facilities. Total power use for the computer room area of these data centers is no more than 40 W/square foot, including all auxiliary power use and cooling energy. There are two draft journal articles from this work, one focusing on the detailed power use of the data center we've examined in most detail, and the other presenting the aggregate electricity use associated with hosting-type data centers in the U.S.
- Mitchell-Jackson, J., J. G. Koomey, B. Nordman, and M. Blazek. 2001. Data Center Power Requirements: Measurements From Silicon Valley. *Energy—the International Journal (Under review)*. <http://enduse.lbl.gov/Projects/InfoTech.html>
 Current estimates of data center power requirements are greatly overstated because they are based on criteria that incorporate oversized, redundant systems, and several safety factors. Furthermore, most estimates assume that data centers are filled to capacity. For the most part, these numbers are unsubstantiated. Although there are many estimates of the amount of electricity consumed by data centers, until this study, there were no publicly available measurements of power use. This paper examines some of the reasons why power requirements at data centers are overstated and adds actual measurements and the analysis of real-world data to the debate over how much energy these facilities use.
- Nordham, Reiss, and Stein. 2001. Delivering Energy Services to Internet Hotels and Other High Density Electronic Loads, Part I: Structure of the HiDEL Industry. Platts Research and Consulting, Boulder, CO.
- Patel, C. D., C. E. Bash, C. Belady, L. Stahl, and D. Sullivan. 2001. Computational Fluid Dynamics Modeling of High Compute Density Data Centers to Assure System Inlet Air Specifications. Reprinted from the proceedings of the Pacific Rim ASME International

Electronic Packaging Technical Conference and Exhibition (IPACK 2001), © 2001, ASME.

Due to high heat loads, designing the air conditioning system in a data center using simple energy balance is no longer adequate. Data center design cannot rely on intuitive design of air distribution. It is necessary to model the airflow and temperature distribution in a data center. This paper presents a computational fluid dynamics model of a prototype data center to make the case for such modeling.

Patel, C. D., R. Sharma, C. E. Bash, and A. Beitelmal. 2002. Thermal Considerations in Cooling Large Scale High Compute Density Data Centers. 8th ITherm Conference. San Diego CA.

A high compute density data center of today is characterized as one consisting of thousands of racks each with multiple computing units. The computing units include multiple microprocessors, each dissipating approximately 250 W of power. The heat dissipation from a rack containing such computing units exceeds 10 KW. Today's data center, with 1000 racks, over 30,000 square feet, requires 10 MW of power for the computing infrastructure. A 100,000 square foot data center of tomorrow will require 50 MW of power for the computing infrastructure. Energy required to dissipate this heat will be an additional 20 MW. A hundred thousand square foot planetary scale data center, with five thousand 10 KW racks, would cost ~\$44 million per year (@ \$100/MWh) just to power the servers & \$18 million per year to power the cooling infrastructure for the data center. Cooling design considerations by virtue of proper layout of racks can yield substantial savings in energy. This paper shows an overview of a data center cooling design and presents the results of a case study where layout change was made by virtue of numerical modeling to avail efficient use of air conditioning resources.

PG&E. 2001. Data Center Energy Characterization Study. Pacific Gas and Electric Company (subcontractor: Rumsey Engineers), San Francisco, Feb. 2001.

Rumsey Engineers, Inc. and PG&E have teamed up to conduct an energy study as part of PG&E's Data Center Energy Characterization Study. This study will allow PG&E and designers to make better decisions about the design and construction of data centers in the near future. Three data centers in the PG&E service territory have been analyzed during December 2000 and January 2001, with the particular aim of determining the end-use of electricity. The electricity use at each facility was monitored for a week each. At the end of the report are a set of definitions, which explain the terms used and the components in making each calculation. The three data centers provide co-location service, which is an unmanaged service that provides rack space and network connectivity via a high capacity backbone. About half or more of the electricity goes to powering the data center floor, and 25 to 34 percent of the electricity goes to the heating, air conditioning and ventilation equipment. The HVAC equipment uses a significant amount of power and is where energy efficiency improvements can be made. All three facilities use computer room air conditioning (CRAC) units, which are stand-alone units that create their own refrigeration and circulate air. A central, water-cooled chilled water system with air handlers and economizers can provide similar services with roughly a 50% reduction in cooling energy consumption. Energy density of the three buildings had an average of 35 W/sf. The cooling equipment energy density for the data center floor alone averaged at 17 W/sf for the three facilities. The average designed energy density of the three data centers' server loads was 63 W/sf, while the measured energy density was 34 W/sf. An extrapolated value was also calculated to determine what the server load energy density would be when fully occupied. The average extrapolated energy density was 45 W/sf. Air movement efficiency varies from 23 to 64 percent between the three facilities. Cooling load density varies from 9 to 70 percent between the three facilities.

Planet-TECH. 2002. Technical and Market Assessment for Premium Power in Haverhill. Planet-TECH Associates for The Massachusetts Technology Collaborative, www.mtpc.org, Westborough, MA 01581-3340, Revision: February 20, 2002.
http://www.mtpc.org/cluster/Haverhill_Report.pdf; <http://www.planet-tech.com/content.htm?cid=2445>

This study is pursued under contract to the Massachusetts Technology Collaborative, in response to a request for a "Technical and Market Assessment". It seeks to determine if the provisioning of "premium

power" suitable for data-intensive industries will improve the marketability of a Historic District mill building in Haverhill. It is concluded that such provisioning does improve the marketability, however, not to a degree that is viable at this time. Other avenues for energy innovation are considered and recommendations for next steps are made.

RLX Technologies. 2001. *Redefining Server Economics*. White Paper.

http://www.rlx.com/pdf/RLXServerEconWP_v1.0.pdf

RMI, and DR International. 2002. Energy Efficient Data Centers - A Rocky Mountain Institute Design Charrette. Organized, Hosted and Facilitated by Rocky Mountain Institute, with D&R International, Ltd. and Friends. Hayes Mansion Conference Center, San Jose, California. <http://www.rmi.org/sitepages/pid626.php>

Rapid growth of "mission critical" server-farm and fiber-optic-node data centers has presented energy service providers with urgent issues. Resulting costs have broad financial and societal implications. While recent economic trends have severely curtailed projected growth, the underlying business remains vital. The current slowdown allows us all some breathing room—an excellent opportunity to step back and carefully evaluate designs in preparation for surviving the slowdown and for the resumption of explosive growth. Future data center development will not occur in the first-to-market, damn-the-cost environment of 1999-2000. Rather, the business will be more cost-competitive, and designs that can deliver major savings in both capital cost (correct sizing) and operating cost (high efficiency)—for both new build and retrofit—will provide their owners and operators with an essential competitive advantage.

Robertson, C., and J. Romm. 2002. Data Centers, Power, and Pollution Prevention - Design for Business and Environmental Advantage. The Center for Energy and Climate Solutions; A Division of The Global Environment and Technology Foundation, June 2002.

<http://www.cool-companies.org>; <http://www.getf.org>

Computers and other electronic equipment will crash at the slightest disruption or fluctuation in their supply of electricity. The power system was not designed for these sensitive electronic loads and is inherently unable to meet the technical requirements of the information economy. For data centers, which play a central role in the information economy, crashing computers cause potentially catastrophic financial losses. The same voltage sag that causes the lights to dim briefly can cause a data center to go off-line, losing large sums of money, for many hours. Data center owners and their power providers must therefore solve several related technical and economic electric power problems. These are: 1) How to assure high-availability (24x7) power supply with a very low probability of failure; 2) How to assure practically perfect power quality; and 3) How to manage risk while minimizing capital and operating expenses

Roth, K. W., Fred Goldstein, and J. Kleinman. 2002. Energy consumption by office and telecommunications equipment in commercial buildings, Volume I: Energy Consumption Baseline. Arthur D. Little (ADL), Inc., 72895-00, Cambridge, MA, January 2002.

ADL carried out a "bottom-up" study to quantify the annual electricity consumption (AEC) of more than thirty (30) types of non-residential office and telecommunications equipment. A preliminary AEC estimate for all equipment types identified eight key equipment categories that received significantly more detailed studied and accounted for almost 90% of the total preliminary AEC. The Key Equipment Categories include: *Computer Monitors and Displays, Personal Computers, Server Computers, Copy Machines, Computer Network Equipment, Telephone Network Equipment, Printers, Uninterruptible Power Supplies (UPSs)*. The literature review did not uncover any prior comprehensive studies of telephone network electricity consumption or uninterruptible power supply (UPS) electricity consumption. The AEC analyses found that the office and telecommunications equipment consumed 97-TWh of electricity in 2000. The report concludes that commercial sector office equipment electricity use in the U.S. is about 3% of all electric power use. The ADL work also creates scenarios of future electricity use for office equipment, including the energy used by telecommunications equipment.

Shields, H. and C. Weschler, 1998. Are Indoor Pollutants Threatening the Reliability of Your Electronic Equipment? Heating/Piping/Air Conditioning Magazine. May.

Stein, Jay. 2002. More Efficient Technology Will Ease the Way for Future Data Centers. Proceedings 2002 ACEEE Summer Study on Energy Efficiency in Buildings.

Sullivan, R. F. 2002. Alternating Cold and Hot Aisles Provides More Reliable Cooling for Server Farms. The Uptime Institute. <http://www.uptimeinstitute.org/tuiaisles.html>
The creation of "server farms" comprising hundreds of individual file servers has become quite commonplace in the new e-commerce economy, while other businesses spawn farms by moving equipment previously in closets or under desktops into a centralized data center environment. However, many of these farms are hastily planned and implemented, as the needed equipment must be quickly installed on a rush schedule. The typical result is a somewhat haphazard layout on the raised floor that can have disastrous consequences due to environmental temperature disparities. Unfortunately, this lack of floor-layout planning is not apparent until after serious reliability problems have already occurred.

The Uptime Institute. 2000. Heat-Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment. The Uptime Institute, Version 1.0.
<http://www.upsite.com/> / <http://www.uptimeinstitute.org/heatdensity.html>
This white paper provides data and best available insights regarding historical and projected trends in power consumption and the resulting heat dissipation in computer and data processing systems (servers and workstations), storage systems (DASD and tape), and central office-type telecommunications equipment. The topics address the special needs of Information Technology professionals, technology space and data center owners, facilities planners, architects, and engineers.

Thompson, C. S. 2002. Integrated Data Center Design in the New Millennium. *Energy User News*.
http://www.energyusernews.com/CDA/ArticleInformation/features/BNP_Features_Item/0,2584,70578,00.html
Data center design requires planning ahead and estimating future electrical needs. Designers must accurately predict space and energy requirements, plus cooling needs for new generations of equipment. Importance of data center reliability is discussed.

Uptime Institute, 2000. Heat Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment. Santa Fe, NM.
<http://www.upsite.com/TUIpages/whitepapers/tuiheat1.0.html>

Wood, L. 2002. Cutting Edge Server Farms - The blade server debate. [newarchitectmag.com](http://www.newarchitectmag.com).
<http://www.newarchitectmag.com/documents/s=2412/na0702f/index.html>. July 23, 2002.
A blade is the industry term for a server that fits on a single circuit board, including CPU, memory, and perhaps a local hard disk. Multiple blades are plugged into a chassis, where each blade shares a common power supply, cooling system, and communications back plane. Multiple chassis can then be stacked into racks. By comparison, the conventional approach for rack-mounted servers involves only one server per chassis. A chassis cannot be smaller than one vertical rack unit (1U, or about 1.75 inches high). This limits you to 42 to 48 servers in a standard seven-foot rack. A typical blade chassis is much higher than 1U, but several can still be stacked in a rack, allowing upwards of 300 servers per rack, depending on the vendor and configuration. This compact design offers compelling advantages to anyone operating a high-density server farm where space is at a premium.

Indeed, blades are the "next big thing" in servers, and it's probable that any given administrator will have to decide whether to adopt them in the near future.

Appendix XIV - RMI Charrette brochure, attendee list, and report.



Low-Power Data Centers

A ROCKY MOUNTAIN INSTITUTE DESIGN CHARRETTE

February 2 to 5, 2003
Hayes Mansion Conference Center, San Jose, California

Sponsors

GENERAL

Pacific Gas & Electric
<http://www.pge.com>

California Energy Commission
<http://www.energy.ca.gov/>

New York State Energy Research and Development Authority (NYSERDA)
<http://www.nyserda.org>

ADEME (Agence de l'Environnement et de la Maîtrise de l'Energie)
Vanbonne, France
<http://www.ademe.fr/>

Los Angeles Dept. of Water & Power (LADWP)
<http://www.ladwp.com>

Southern California Edison Company
<http://www.edison.com>

BC Hydro Power Smart
<http://www.bchydro.com>

The Canton of Geneva (Switzerland)

- Département de l'intérieur, de l'agriculture et de l'environnement (DIAE)
<http://www.geneve.ch/diae/welcome.asp>
- Service cantonal de l'énergie (ScanE)
<http://www.geneve.ch/scane/home/welcome.asp>

Centre for Energy Policy & Economics (CEPE)
Zurich, Switzerland
<http://www.cepe.ethz.ch>

California Independent System Operators (ISO)
<http://www.caiso.com>

INDUSTRY / PARTICIPANT SUPPORT

American Power Conversion
<http://www.apcc.com>

RealEnergy, Inc.
<http://www.realenergy.com>

S&C Electric Company
Power Quality Products Division
<http://www.sandc.com>

Wuchun Feng
Los Alamos National Lab

EYP Mission Critical Facilities, Inc.
<http://www.eypae.com>

Glumac International
<http://www.Glumac.com>

Jones Lang LaSalle Americas, Inc.
<http://www.joneslanglasalle.com>

Engineering Enterprise
<http://www.engent.com>

IN KIND

Lawrence Berkeley National Lab (LBNL)

John Gage, Sun Microsystems
<http://www.johngage.com>

Oregon Office of Energy
<http://www.state.or.us>

The Natural Step
<http://www.naturalstep.org>

Natural Resources Defense Council (NRDC)
<http://www.nrdc.org>

RMI Staff and Consultant Team

Rocky Mountain Institute

1739 Snowmass Creek Road
Snowmass, CO 81654
970-927-3851

William Browning, principal and founder of RMI's
Green Development Services
970.927.7304 / browning@rmi.org

Cameron Burns, RMI writer and editor
970.927.7338 / cameron@rmi.org

Odd-Even Bustnes, special aide to Amory Lovins
970.927.7343 / oebustnes@rmi.org

Craig Collins, RMI Energy & Resources Services
craigcollins@rmi.org

Huston Eubank, AIA, principal and charrette co-
organizer, a member of RMI's Green Development
Services 970.927.7317 / huston@rmi.org

Marty Hagen, RMI information systems manager
970.927.7314 / mhagen@rmi.org

Onno Koelman, RMI's Mineral Acquisitions Partnership
Fellowship intern
970.927.7301 / okoelman@rmi.org

Corey Griffin, Former RMI/GDS Konheim Intern
cgriffin@uclink.berkeley.edu

Amory Lovins, Convener, chief executive officer of
Rocky Mountain Institute
970.927.3851 / ablovins@rmi.org

Jenifer Seal, principal, a member of RMI's Green
Development Services team
970.927.7322 / jenseal@rmi.org

Joel Swisher, PhD, PE, principal and team leader of
RMI's Energy & Resources Services 303.245.1003 /
jswisher@rmi.org

Tom Watanabe, charrette co-organizer, RMI business
development director
510.526.5644 / tomwatanabe@rmi.org

Devra Bachrach

Energy Scientist
Natural Resources Defense Council (NRDC)
71 Stevenson St., #1825
San Francisco, CA 94105

415-777-0220

dbachrach@nrdc.org

Dick Bourne

Davis Energy Group Inc.
123 C St
Davis, CA 95616
530.753.1100
dbourne@davisenergy.com

Malcolm Lewis

Constructive Technologies Group (CTG)
16 Technology Drive, Suite 109
Irvine, CA 92618
949-790-0010
mlewis@ctg-net.com

Piers Heath

Battle McCarthy
Dog and Duck Yard
Princeton Street
London WC1R 4BH
+44 (0)20 7 440 8261
pheath@battlemccarthy.com

Gregory Kats

Principal
Capital E
(Former DOE, author of IPMVP)
1828 L. Street NW
Suite 1046
Washington, DC 20003
202-463-8469
<http://www.cap-e.com/>
gkats@cap-e.com

Ron Perkins

Supersymmetry USA Inc
4509 Evergreen Forest Lane
Navasota, TX 77868
936-894-2819
ronaldperkins@msn.com

Peter Rumsey

Rumsey Engineers Inc
99 Linden St
Oakland, CA 94607-2508
510-663-2070
prumsey@rumseyengineers.com

Dale Sartor

Environmental Energy Technologies Division
Building Technologies Department
Lawrence Berkeley Nat. Lab (LBNL)
MS 90- 3136
Berkeley, CA 94720
510-486-5988

<http://eetd.lbl.gov/btp/ateam/ateamweb/team/dale.htm>
DA_Sartor@lbl.gov

JB Straubel
Volacom
1509 East Garfield Ave.
Glendale, CA 91205
650-996-4355
straubel@stanfordalumni.org

Geoff Wood
Profile Composites Inc
1416 Lands End Road
Sidney, BC V8L 5K1
250-655-7142
GMWood@aol.com

Gautam Barua
The Natural Step
116 New Montgomery Street
Suite 800
San Francisco, CA 94105
415-531-5317
gbarua@naturalstep.org

Michael Daish
616 Pomeroy Avenue
Santa Clara Ca 95051
408 455 3145
mdaish@email.com

Greg McBeth, Stanford student
piguy@Stanford.EDU

Cody Taylor, Stanford student
PO Box 13698
Stanford CA 94309
cody1984@stanford.edu

Participants

Barry Abramson
Sr. Vice President, Engineering
Servidyne Systems
LLC 1350 Spring Street, NW, Suite 200
Atlanta, GA 30309
404-352-2050 X 255
barry.abramson@servidyne.com

Eric Adrian
Jones Lang LaSalle
200 East Randolph
Chicago, IL 60601
eric.adrian@am.joneslanglasalle.com

Bernard Aebischer
Centre for Energy Policy & Economics (CEPE)
ETH Zentrum, WEC
8092 Zurich
+41-1-632 41 95
www.cepe.ethz.ch
bernard.aebischer@cepe.mavt.ethz.ch

Adrian Altenberg
Anstein +Walther AG
Leutschenbachstrasse 45, 8050
Zürich Switzerland
+41 305 91 11
www.amstein-walthert.ch
adrian.altenburger@amstein-walthert.ch

Alain Anglade
ADEME (Agence de l'Environnement et de la Maîtrise
de l'Energie)
Centre de Sophia Antipolis
500 route des Lucioles,
06500 Vanbonne, France
011 31 4 9395 7935
alain.anglade@ademe.fr

Dan Baer
VP, Environmental Products
Member of Telcordia Comm. on Thermal Mgmt in DCs
Liebert Corporation
1050 Dearborn Drive
Columbus, OH 43085-1544
614-841-5723
<http://www.Liebert.com>
<http://www.cha-portal.com/>
Dan.Baer@liebert.com

Kevin Best
Director of Sales & Marketing
RealEnergy, Inc.
300 Capitol Mall, Suite 120
Sacramento, CA 95814
916-325-2500 X 109
707-426-3300
kbest@realenergy.com

Clark Bisel
Senior Vice President
Flack + Kurtz Consulting Engineers
343 Sansome Street, Suite 450
San Francisco, CA 94104
415-402-2231
Clark.Bisel@sf.fk.com

Kenneth Brill
Executive Director
The Uptime Institute, Inc.
1347 Tano Ridge Road

Santa Fe, NM 87506
505-982-0514
505-577-2098 cell
<http://www.upsite.com>
tui@upsite.com

Chris Chouteau
Senior Vice President
Pacific Gas & Electric
PO Box 770000
San Francisco, CA 94177
415-972-5239
wcc2@pge.com

Tom Coulard
Data Support Associates, Inc.
P.O. Box 220
Brookfield, CT 06804
203.740.4222
516.236.9198
tcoulard@datasupportassociates.com

Dave Coup
Project Manager
New York State Energy Research and Development
Authority (NYSERDA)
17 Columbia Circle
Albany, NY 12203-6399
518-862-1090 X 3379
jdc@nyserda.org

Tom Croda
Principal Engineer
ClearBlue Technologies
100 First Street, Suite 2000
San Francisco, CA 94105
415-869-7100
tom.croda@clearblue.com

Joe Daniels
Engineer, Facilities Group, Power Section
Department of Defense / National Security Agency
301-688-5560
joedaniels711@netscape.net

Tom Ditoro
HDR Architects
8404 Indian Hills Drive
Omaha, NE 68144
402-399-1169
<http://www.hdrinc.com>
tditoro@hdrinc.com

Patsy Dugger
Program Manager
Industrial and Agricultural New Construction
Customer Energy Management
Savings By Design Program

Pacific Gas and Electric Company
PO Box 770000
San Francisco, CA 94177
415-973-1019
pwd2@pge.com

Grant Duhon
Program Manager Supervisor
Customer Energy Management
Savings By Design Program
Pacific Gas and Electric Company
Mail Code: N6G
PO Box 770000
San Francisco, CA 94177
415-972-5695
gxd7@pge.com

Gary Engle
California ISO
151 Blue Ravine
Folsom CA 95630
916-608-5974
gengle@caiso.com

Wu-chun Feng
Los Alamos National Lab
CCS-1 MS-D451
LANL
Los Alamos, NM 87545
505-665-2730
feng@lanl.gov

Stephen Fok
Senior Program Engineer
Technical Application Services
Customer Energy Management
Savings By Design Program
Pacific Gas and Electric Company
415-973-4735
SKF2@pge.com

Rafael Friedman
Senior Project Manager
Customer Energy Management
Mail Code N6G
Pacific Gas & Electric Company
P.O.Box 770000
San Francisco, CA 94177-0001
415-972 5799
415-972-4799
RAFi@pge.com

John Gage
Chief Scientist
Sun Microsystems
P.O. Box 2065
Burlingame, CA 94011
650-786-3400

Ass't: Tricia Fitzgerald
patricia.fitzgerald@sun.com
www.johngage.com
john.gage@sun.com

Angelina Galiteva
Los Angeles Dept. of Water & Power (LADWP)
111 N. Hope Street
Room 1551
Los Angeles, CA 90012
213-367-1577
angelina.galiteva@ladwp.com

Steve Greenberg
President
Thin Client Computing
34522 N. Scottsdale Rd. D8453
Scottsdale, AZ 85282
602-432-8649
<http://www.thinclient.net>
steveg@thinclient.net

Steven Greenberg
Senior Energy Officer
RealEnergy, Inc.
5957 Variel Avenue
Woodland Hills, CA 91367
818-610-2300
sgreenberg@realenergy.com

Peter Gross
CEO
EYP Mission Critical Facilities Inc.
11845 W. Olympic Blvd.
Suite 850 West
Los Angeles, CA 90064
310-914-3442
<http://www.eypae.com/>
pgross@eypmcf.com

Don Hall
Account Services Supervisor
Pacific Gas and Electric Company
111 Almaden Blvd
San Jose, CA
408-282-7444
DJH5@pge.com

Chris Hipp
Founder and former CTO
RLX Technologies
founding member
Server Blade Trade Association
1590 Stockbridge Ave.
Redwood City, CA 94061-4222
650-455-8881
chris@hippster.com

Ron Hughes
Principal
California Data Center Design Group
14925 Lago Drive
Rancho Murieta, Ca. 95683
916-531-7396
rhughes@cdcdg.com

Ann Hushagen
Energy Analyst
Oregon Office of Energy
625 Marion St. N.E.
Salem, OR 97301
503-373-7804
Ann.Hushagen@state.or.us

Jerry Hutchinson
Major Account Rep - Data Centers
Account Services
Pacific Gas and Electric Company
111 Almaden Blvd
San Jose, CA
408-282-7345
JAHo@pge.com

Kevin Hydes
President & CEO
Keen Engineering Co Ltd
116 - 930 West 1st Street
North Vancouver, BC
V7P 3N4 Canada
604-986-5336
kevin.hydes@keen.ca

Ernie Jensen
Utility Manager
Telus Communications
12040-107 Street
Edmonton, Alberta, Canada T5G 2S7
780-493-7610
ernie.jensen@telus.com
<http://www.telus.com>

Steve Jurvetson
Managing Director
Draper Fisher Jurvetson
400 Seaport Court, Suite 250
Redwood City, CA 94063
650-599-9000
<http://www.drapervc.com/>
SJ@dfj.com

Ron Kalich
Director, Facilities
ClearBlue Technologies
100 First Street, Suite 2000
San Francisco, CA 94105

415-869-7116
ron.kalich@clearblue.com

Mukesh Khattar

Energy Director
Oracle Corporation
500 Oracle Parkway, MS LGN-2
Redwood Shores, CA 94065
650-506-6980
mukesh.khattar@oracle.com

Jonathan Koomey

Staff Scientist, End-use Forecasting Group
Lawrence Berkeley Nat. Lab (LBNL)
1 Cyclotron Road, Building 90-4000
Berkeley, CA 94720
510-486-5974
<http://enduse.lbl.gov>
JGKoomey@lbl.gov

Henry Lau

P.E., Design and Engineering Services
Southern California Edison
6042 N. Irwindale Ave., Ste. B
Irwindale, CA 91702
626-633-7179
lahu@sce.com

Donald Lee

Senior Vice President
Bank of America
#10609
CA5-701-10-07
1455 Market St., 10th Floor
San Francisco, CA 94103
donald.t.lee@bankofamerica.com

Jim Magdych

CIO
Cool Chips plc
408-621-6125
jmagdych@coolchips.gi

KC Mares

Director, Data Center Operations
Redundant Networks
9408 Double R Blvd., Suite B
Reno, NV 89521
775-850-4222 X 2237
<http://www.redundant.com>
kcmars@redundant.com

Greg Mears

VP, Engineering
Invensys Powerware
8609 Six Forks Rd
Raleigh, NC 27615

919-878-1003
greg.mears@invensys.com

Ben Mehta

Major Account Rep - Cleanrooms
Account Services
Pacific Gas and Electric Company
111 Almaden Blvd
San Jose, CA
408-725-2062
DJH5@pge.com

Bruce Nordman

Lawrence Berkeley Nat. Lab (LBNL)
1 Cyclotron Road
90-4000
Berkeley, CA 94720-8136
510-486-7089
BNordman@lbl.gov

John Pappas

Principal
Mazzetti and Associates
530 Bush Street, Third Floor
San Francisco, CA USA 94108
415-362-3266
www.mazzetti.com
johnp@mazzetti.com

Chandrakant Patel

P.E., Principal Scientist
Internet Systems and Storage Laboratory
Hewlett-Packard Laboratories
1501 Page Mill Rd
Palo Alto, CA94304
650-857-7140
650-743-1599 cell
chandrakant_patel@hp.com

Bob Perreault

Caliber Facilities Mgt. Ltd.
1801 Broadview Rd NW
Calgary, AB Canada T2N 3H6
403-270-3223
calfac@telusplanet.net

Neil Rasmussen

SVP, Chief Technical Officer
American Power Conversion
85 Rangeway Road, Bldg 2
North Billerica, MA 01862
<http://www.apcc.com>
neil.rasmussen@apcc.com

Brad Roberts

S&C Electric Co.
Power Quality Products Division
5251 West Franklin Drive

Franklin, WI 53132
414-423-8776 X 109
broberts@sandc.com

Paul Roggensack
Mechanical Engineer
California Energy Commission
1516 Ninth Street, MS 43
Sacramento, CA 95814-5504
916-654-6560
Proggens@energy.state.ca.us

Art Rosenfeld
Commissioner
California Energy Commission
1516 9th Street, MS-35
Sacramento, CA 95814
916-654-4930,
Arosenfe@energy.state.ca.us

David Schirmacher
Goldman Sachs
180 Maiden Lane
15th Floor
New York, NY, 10038
212.357.7482
david.schirmacher@gs.com

Roland Schoettle
Chief Executive Officer
Optimal Technologies International Inc.
1513 Spyglass Parkway
Valljo, CA 94591
707-557-1788
916-718-9274 cell
rolands@otii.com

Steve Schumer
VP, Technology
Cupertino Electric Energy Group (CEI Energy Group)
1132 North Seventh St.
San Jose, CA 95112
408-808-8173
408-966-9550 cell
Steve_Schumer@cei.com

Bob Seese
Manager, Data Center Services
Loudcloud
599 N. Mathilda
Sunnyvale, CA 94085
408-212-4630
<http://www.loudcloud.com>
bseese@loudcloud.com

Peter Spark
Principal
Ecsponent

RO9022
Flat 8, 37-39
Great Ormond Street
London WC1N3HZ
pspark.ocs2000@london.edu

Mike Steinmann
General Manager, San Francisco
Glumac International
10419 Old Placerville Rd # 250
San Francisco, CA 94107
415-398-7667
msteinmann@glumac.com
www.Glumac.com

Joe Stolarski
SVP, Director of Engineering and Operations
Jones Lang LaSalle
200 East Randolph
Chicago, IL 60601
312-228-2091
joe.stolarski@am.joneslanglasalle.com

Steve Strauss
President
Glumac International
320 SW Washington, Suite 200
Portland, Oregon 97204
503-227-5280
www.Glumac.com
sstrauss@glumac.com

Harold J. Stewart
Principal
Keen Engineering Co. Ltd.
116 - 930 West 1st Street
North Vancouver, BC
Canada V7P 3N4
604-986-5336
604-983-4034 direct
604-219-8385 cell
harold.stewart@keen.ca

Steven Torres
VP, Western Region
Fuel Cell Energy, Inc.
225 South Lake Avenue, Suite 300
Pasadena, CA 91101
626-432-5410
626-818-4647 cell
storres@fce.com

William L True
FCRE Corporate Engineering
82 Devonshire Street, W3B
Boston, MA 02109
617-563-7349

bill.true@fmr.com

William Tschudi

Lawrence Berkeley Nat. Lab (LBNL)
1 Cyclotron Rd. Ms 90R3111
Berkeley, CA 94720
510-486-5988
wtschudi@lbl.gov

Ashlee Vance

IDG News Service
San Francisco Bureau
San Francisco, CA 94107
415-974-7391
Ashlee_Vance@idg.com

James Warren

Atlantic Regional Director
EYP Mission Critical Facilities, Inc.
54 State Street, 3rd Floor
Albany, NY 12207
518-337-2013
jwarren@eypmcf.com

Bill Westbrook

Business Development Manager
Power Measurement USA
1099 D St Suite 208
San Rafael CA 94901
877-476-5872
415-457-9040 direct
415-595-6716 mobile
bill_westbrook@pml.com

Scott Wheeler

Principal
The Engineering Enterprise Consulting Engineers
853 Lincoln Way, Suite 109
Auburn, CA 95603
530-886-8556
tee-sw@pacbell.net
<http://www.engent.com>

Richard Williams

Jones Lang LaSalle
200 East Randolph
Chicago, IL 60601
richard.williams@am.joneslanglasalle.com

John Wilson

Special Ass't to Art Rosenfeld
California Energy Commission
1516-9th St.
Sacramento, CA 95814-5512
916-654-5056
jwilson@energy.state.ca.us

Ron Wilson

Office Director
M&A SiliconValley
452 South First Street
San Jose, CA 95113
408-279-2390
RonW@mazzetti.com

Tim Xu

Project Manager, Building Technologies Department
Environmental Energy Technologies Division
Lawrence Berkeley Nat. Lab (LBNL)
One Cyclotron Road
BLDG 90R3111
Berkeley, CA 94720-8134
510-486-7810
ttxu@lbl.gov

Design Recommendations for High-Performance Data Centers



Report of the
Integrated Design Charrette
Conducted
2-5 February 2003





This document may be reprinted with proper credit
to Rocky Mountain Institute (RMI).

Please visit www.rmi.org/sitepages/pid626.php
for the report, addenda, and ongoing corrections and updates.

This report was written, edited, and designed by
Huston Eubank, Joel Swisher, Cameron Burns, Jen Seal,
and Ben Emerson, with funding from
Pacific Gas & Electric Energy Design Resources program.
Hardware photos throughout the report are courtesy
Chris Hipp, charrette photos are by Cameron Burns.

RMI Published by Rocky Mountain Institute
1739 Snowmass Creek Road
Snowmass, CO 81654-9199, USA

phone: 1.970.927.3851
fax: 1.970.927.4178

www.rmi.org

Cover: 2001 RLX Technologies™ ServerBlade 667—originally designed
by RLX founder Chris Hipp.

Type: Palatino (body text) and Univers (supporting text elements).
Paper: New Leaf™ Reincarnation Matte 95#, 100% recycled, 50% PCW, PCF (cover).
and New Leaf™ Eco Offset 60#, 100% PCW, PCF (text).
Printed with vegetable-based ink and bound with biodegradable glue

Printed in the United States of America

© 2003 Rocky Mountain Institute (RMI) or other copyright holders.
The trademarks contained in this report are owned and/or controlled by
Rocky Mountain Institute or by other parties and are the property of
their respective owners. It is possible that some elements of this report
may be the subject of patent rights. Rocky Mountain Institute shall not be held
responsible for identifying any such patent rights.

Whether, how, and to what extent the recommendations in this report
are properly or fully implemented is at the discretion of individual users of
this report. Therefore RMI cannot be held liable for any consequences or
damages that may arise from implementation or non-implementation
of these recommendations or from reliance on
any information contained in this report.

Appendices

- A** Wu-chun Feng: *Bladed Beowulf: Cost Effective Alternatives to Traditional Beowulfs*
- B** Panel Discussion: "Issues in the Design of Data Centers"
- C** Jon Koomey/Bill Tschudi Presentation
- D** Peter Rumsey Presentation
- E** Chris Hipp Presentation: *Blade Computing: How Did We Get Here? Where Are We Going?*
MS Word doc (1), MS Power Point (2)
- F** Piers Heath Presentation
- G** Amory Lovins Presentation
- H** B. Aebischer and A. Huser, *Energy Efficiency of Computer Power Supplies*,
(abstract in English) November 2002
- I** NRDC Report: *Power Supplies: A Hidden Opportunity for Energy Savings*
- J** *Thermal Considerations in Cooling Large Scale High Compute Density Data Centers*,
by Chandrakant Patel *et al.* (1); and *Computational Fluid Dynamics Modeling of High Compute Density
Data Centers* by Chandrakant Patel *et al.* (2)
- K** *Cool Chips Overview*, by Jim Magdych of Cool Chips, Inc.
- L** Notes from Heat Sources and Transfer (breakout) Group
- M** *Data Center Energy Characterization Study*, Rumsey Engineers, February 2001
- N** "N1's Computing-on-Demand to Drive Network Services" by Greg Papadopolous,
Network Computing Asia, 1 February 2003
- O** "Transmeta Announces Features of Next Generation TM8000 Processor for Energy Efficient
Computing," Transmeta press release, 10 March 2003
- P** "Dockable Server Concepts," Will Berry and Stephen W. Montgomery, Intel Labs, 25 February 2003
- Q** "Powering the Internet, Datacom Equipment in Telecom Facilities: The Need for a DC Powering Option,"
by the Technical Subgroup on Telecommunications Energy Systems of the Power Electronics Society
of the Institute of Electrical and Electronics Engineers, Inc., 1998
- R** Energy Performance Contracting for New Buildings
- S** *Energy- and Eco-Efficiency of Data Centres: A study commissioned by DIAE1 / ScanE2
of the Canton of Geneva*, by B. Aebischer *et al.*, 5 January 5 2003
- T** "DARPA funds power-aware architecture development" by Stephan Ohr, *EE Times*, 17 August 2003
- U** "Cooling and Power Considerations for Semiconductors Into the Next Century," by Christian Belady

Design Recommendations for
High-Performance
Data Centers

5	Acknowledgments
9	Introduction
14	Summary
25	Matrix of Recommendations
32	Recommendations
83	Conclusion

85	About Rocky Mountain Institute (RMI)
86	RMI Team Directory and Bios
90	Participant List and Bios
100	Sponsors

<< **Appendices**

*Table of Appendices accompanies the CD-ROM on inside front cover
Appendices are furnished in electronic format only.
Contact orders@rmi.org for additional copies.*

Part 1: Native Loads: CPUs, Servers, Software, etc.

Part 2: Computer Power Supplies


Part 3: Next Generation Cooling

Part 4: Cooling

Part 5: Facility Power Supply

Part 6: Operations


*Parts 1–6 include 62 total recommendations
outlined on next page*



Recommendations

A three-day transdisciplinary workshop with ~90 industry experts synthesized ways to design, build, and operate data centers that would use approximately tenfold less energy, improve uptime, reduce capital cost, speed construction, and enhance the value proposition. Some of the >50 integrated recommendations can also be applied to existing data centers.

Part 1: Native Loads (CPUs, Servers, Software, etc.) 33	Part 4 Cooling 51	Part 5: Facility Power Supply 64
.1 New design paradigm 34	a: Low Energy Ventilation and Cooling 52	.1 AC power distribution system 66
.2 Define operating envelope 36	.1 Increase the temperature range of cooling air 52	.2 On-site power generation 70
.3 Reduce or eliminate heat sources 37	.2 Manage airflow to reduce energy required for cooling and ventilation 52	.3 Interconnect with utility 71
.4 High efficiency CPUs 37	.3 Minimize air-side static pressure 54	.4 Address barriers to self-generation 71
.5 Remove disk drives from servers 38	.4 Maximize use of free cooling 55	
.6 Power supplies 38	.5 Natural ventilation 56	Part 6: Operations 73
.7 Remove power supplies from servers 39	.6 Demand-controlled ventilation 56	.1 Intelligent resource allocation 75
.8 Dynamically allocate resources 39	.7 Additional ideas 57	.2 Improve information available 75
.9 Create an Energy Star standard for servers 40	.8 Wish list for manufacturers 57	.3 Align incentives with desired performance 76
Part 2: Computer Power Supplies 41	b: Efficient Heat Rejection in Large Data Centers 58	.4 Benchmarking 76
.1 Recognize and account for the full cost of each watt of power delivered to the server 43	.1 Tune humidification and dehumidification cycles on existing systems 58	.5 Write more efficient code 77
.2 Create clear connections between power supply design, system efficiency, and power cost, with incentives to support efficient solutions 43	.2 Evaporative condensers/cooling 58	.6 Submetering 77
.3 Focus on finding continuous, not intermittent, power savings 45	.3 Design and install chilled water systems greater than 200 tons to operate at a total of 0.62 kW per ton 58	.7 Measurement and verification (M&V) 77
.4 Establish industry standards to increase power supply efficiency 45	.4 Design and install chilled water systems greater than 60 and less than 200 tons to operate at a total of 0.83 kW per ton 59	.8 Continuous commissioning 78
.5 Improve power supply design 45	.5 Microclimate-specific recommendations for northern U.S. and cold climates 59	.9 Create self-diagnosing/healing systems 78
Part 3: Next Generation Cooling 47	.6 Use waste heat from on-site cogeneration to drive HVAC system 59	.10 Virtual servers 78
.1 Create system standards 49	.7 Desiccant cooling 60	.11 Optimization tools 78
.2 General 49	.8 Thermal storage 60	.12 Miscellaneous 78
.3 Hybrid approaches for near term 49	.9 Wish list for manufacturers 60	.13 Education, outreach and training 79
.4 Conductive thermal path to liquid for future 50	c: Control Strategies 61	.14 Demonstrations 79
	.1 General low-/no-cost optimizations 61	.15 Energy Star and LEED ratings 80
	.2 Establish environmental standards for mechanical and electrical systems by room type, and control to least energy-intensive values 61	.16 Create an independent organization to provide testing, experimentation, education, and demonstrations 81
	.3 Low-/no-cost solutions: CRAC optimization 62	
	.4 Low-/no-cost solutions: reconfigure controls on central air handlers 62	
	.5 Low-/no-cost solutions: reconfigure controls on central plants 62	
	.6 Mid- to high-cost solutions 62	
	.7 Future control systems 63	



Our sponsors, to whom we at RMI are deeply grateful, include:

- Pacific Gas & Electric Company
(www.pge.com)



- California Energy Commission
(www.energy.ca.gov)



- New York State Energy Research and Development Authority (NYSERDA)
(www.nyserda.org)



- Agence de l'Environnement et de la Maîtrise de l'Énergie (ADEME)
(www.ademe.fr)



- Southern California Edison Company
(www.edison.com)



- BC Hydro Power Smart
(www.bchydro.com)



- Canton de Genève (Switzerland), Département de l'Intérieur, de l'Agriculture et de l'Environnement (DIAE)
(www.geneve.ch/diae/welcome.asp)



- The Service Cantonal de l'Énergie (ScanE)
(www.geneve.ch/scane/home/welcome.asp)



- Center for Energy Policy & Economics (CEPE), Zurich, Switzerland
(www.cepe.ethz.ch)



- California Independent System Operator Corporation (ISO)
(www.caiso.com)



This event was organized, facilitated, and partially funded by Rocky Mountain Institute (RMI). Initial organizational efforts were funded by RMI. Subsequently, many sponsors made it possible to hold this unique event. Without the tireless fundraising efforts of Tom Watanabe, charrette co-organizer and business development consultant to RMI, this charrette would not have taken place.



**In addition to our sponsors,
the charrette received significant industry/participant support from:**

- American Power Conversion
(www.apcc.com)



- RealEnergy, Inc.
(www.realenergy.com)



- S&C Electric Company's Power Quality Products Division
(www.sandc.com)



- EYP Mission Critical Facilities, Inc.
(www.eypae.com)



- Glumac International
(www.glumac.com)



- Jones Lang LaSalle Americas, Inc.
(www.joneslanglasalle.com)



- The Engineering Enterprise
(www.engent.com)



The Engineering Enterprise

A full directory of sponsors is included on the back cover of this report.

**To open the charrette RMI organized a panel discussion
on "Issues in the Design of Data Centers."**

**Thanks to our panelists for providing a lively, informed,
and enlightening discussion:**

Dan Baer—VP, Environmental Products, Leibert;
member of Telcordia Committee on Thermal
Management in DCs; member of European
Colocation & Hosting Association;

Ken Brill—Executive Director, Uptime Institute;

Grant Duhon—Savings By Design Program,
Pacific Gas and Electric Company;

Chris Hipp—Founder and former CTO,
RLX Technologies;

KC Mares—Director, Data Center Operations,
Redundant Networks; former co-chair SVMG Energy
Committee; former Director of Electrical Energy,
Utilities, and Special Projects at Exodus;

Chandrakant Patel—Principal Scientist,
Internet Systems and Storage Laboratory, Hewlett-
Packard Laboratories;

Joe Stolarski—SVP, Director of Engineering &
Operations, Jones Lang LaSalle; and

Jim Warren—EYP Mission Critical Facilities, Inc.



A broad range of industry experts accepted the challenge and participated in the charrette, bringing an unusually wide variety of perspectives, tremendous enthusiasm, and refreshing willingness to go beyond conventional thinking. A full directory of participants is included on p. 86. RMI contracted with other specialist consultants listed in Appendix D to provide facilitation of the discussions and special expertise in integrated, whole-systems design. The facilitators and participants in each breakout group were crucial to the cultivation of innovative ideas and out-of-the-box thinking. We would like to recognize and thank them all here.

Native Loads Group

Facilitator: Devra Bachrach.
Recorder: Onno Koelman.
Participants: Clark Bisel, Marty Hagen, Jim Magdych, Peter Rumsey, Bob Seese, Geoff Wood.
Later: Neil Rasmussen and Tom Croda.

Computer Power Supplies Group

Recorder: Greg McBeth.
Participants: Bernard Aebischer, Tom Croda, Neil Rasmussen, J.B. Straubel.

Cooling Group

Facilitator: Bill Browning.
Recorders: Ann Hushagen, Corey Griffin, and Joanie Henderson.
Participants: Barry Abramson, Dan Baer, Dick Bourne, David Coup, Piers Heath, Ernie Jensen, K.C. Mares, Henry Lau, John Pappas, Ron Perkins, Neil Rasmussen, Bill True, Bill Tschudi, Tim Xu, Malcolm Lewis.

Facility Power Supply Group

Facilitators: Odd-Even Bustnes and Joel Swisher.
Recorders: Craig Collins and Greg McBeth.
Participants: Bernard Aebischer, Tom Croda, Michael Daish, Joe Daniels, Tom Ditoro, Gary Engle, Steven Greenberg, Joe Griffith, Peter Gross, Greg Mears, Peter Mikhail, Neil Rasmussen, Brad Roberts, Art Rosenfeld, Mike Steinman, J.B. Straubel, Stephen Torres, James Warren, Bill Westbrook, Scott Wheeler, Ron Wilson.

Operations Group

Facilitators: Greg Kats, Dale Sartor.
Recorders: Gautam Barua and Cody Taylor.
Participants: Eric Adrian, Kevin Best, Ken Brill, Patsy Dugger, Steve Greenberg, Peter Gross, Ron Hughes, Peter Spark, Ron Kalich, Jon Koomey, Bob Perrault, Jen Seal, Steve Strauss, David Schirmacher, Joe Stolarski, Tom Watanabe.

Many participants were not tied to any specific breakout group:

Adrian Altenberg, Chris Chouteau, Tom Coulard, Grant Duhon, Huston Eubank, Stephen Fok, Rafael Friedman, Jerry Hutchinson, Steve Jurvetson, Mukesh Khattar, Donald Lee, Amory Lovins, Ben Mehta, Bruce Nordman, Paul Roggensack, Roland Schoettle, Steve Schumer, Harold J. Stewart, Richard Williams, and John Wilson.

Several of our facilitators and recorders volunteered their time to support this effort.

We would like to thank them and the organizations that support them:

Greg McBeth and Cody Taylor, currently undergraduate students at Stanford; J.B. Straubel of Volacom;¹ Ann Hushagen of the Oregon Office of Energy; Corey Griffin, a former RMI Intern and currently a graduate student at U.C. Berkeley; Devra Bachrach of the Natural Resources Defense Council (NRDC); Dale Sartor of Lawrence Berkeley National Laboratory; and Gautam Barua of The Natural Step.

¹ A new company that is building extremely long-endurance high altitude aircraft using completely composite air-frames and completely electric propulsion systems that use liquid hydrogen as a fuel.

The following speakers deserve special thanks:

John Gage, Chief Scientist, Sun Microsystems, for his overview of issues facing the industry and creative approaches to resolving them.

Wu-chun Feng, for setting the tone for the entire charrette with his keynote presentation on the Green Destiny computer at Los Alamos National Laboratory.

Chris Hipp, for his presentation on blade servers and innovative solutions to equipment requirements in data centers.

Peter Rumsey, for his presentation on “Energy Efficiency Strategies for Data Center HVAC Systems.”

Piers Heath, for his presentation on “Data Centres: A Design Approach and Methods of Climate Control.”

Bernard Aebischer of the Centre for Energy Policy and Economics (CEPE), Swiss Federal Institute of Technology (ETH), for his presentation on “Energy- and Eco-Efficiency of Data Centres: Past Activities and Future Plans in Geneva.”

Chandrakant Patel, for his presentation on the work of Hewlett-Packard Laboratories on “Smart Data Centers.”

Jon Koomey and **Bill Tschudi** of Lawrence Berkeley National Laboratory (LBNL), for their review of available data from field measurements and utility bills on how much electricity is actually used by data centers, and with what end-use structure.

Jim Magdych, Chief Information Officer, Cool Chips PLC, for his presentation of the development of computer chips using thermotunneling technology.

Dale Sartor of LBNL, who led a special discussion about “next steps.”


Special thanks to California Energy Commissioner Art Rosenfeld for refocusing our attention and efforts early on the important role of equipment power supplies in the overall equation. His intuition was correct.

Also special thanks to Ken Brill, Executive Director of the Uptime Institute, for his energetic participation and valuable insights.

RMI staff wrote this charrette summary report, which has been reviewed by all participants prior to publication. RMI writers include Huston Eubank, Joel Swisher, Cameron Burns, and Jen Seal. Amory Lovins and Jonathan Koomey provided special technical reviews. Layout and graphics are by Ben Emerson of RMI, unless otherwise noted. Many charrette participants were especially helpful in reviewing and interpreting this report: thanks to Chris Hipp, Tom Croda, Ron Perkins, Geoff Wood, Neil Rasmussen, Bernard Aebischer, J.B. Straubel, Bill True, Brad Roberts, Greg Kats, Ken Brill, Dick Bourne, Jon Koomey, Dale Sartor, Tim Xu, Clark Bisel, Dave Coup, John Pappas, and Will Clift. Onno Koelman’s notes on the Native Loads breakout group, included in Appendix L, are stellar.

Any remaining errors are the authors’ responsibility and should kindly be notified to huston@rmi.org, as should further suggestions and achievements. Errata and addenda will be posted from time to time at www.rmi.org/sitepages/pid626.php.

Finally, project manager Huston Eubank would like especially to acknowledge RMI cofounder and CEO Amory Lovins for his presentations, his determination and optimism, and his visionary leadership, all of which made this event and our continued efforts possible.



Rocky Mountain Institute

1739 Snowmass Creek Road
Snowmass, CO 81654

970-927-3851

970-927-4510 fax

www.rmi.org

“While performance and price/performance are important metrics... the key metrics of this decade will be efficiency, reliability, and availability.”

Wu-chun Feng
 Technical Leader, the RADIANT Group
 Los Alamos National Laboratory


Introduction

The world has become dependent upon reliable data manipulation, storage, retrieval, exchange, and safe-keeping. We use data for most of our modern telecommunications, commercial, financial, national security, military, academic, and government systems, among other things. The central issue regarding data is the reliability of the systems housing and powering them.

In 1965 Gordon Moore, cofounder of Intel, found that “the number of transistors per square inch on integrated circuits had doubled every year since the integrated circuit was invented.” Moore predicted that this trend would continue for the foreseeable future. In subsequent years, the pace slowed down a bit, but data density has doubled approximately every 18 months, and this is the current definition of Moore’s Law, which Moore himself has blessed. Most experts, including Moore himself, expect Moore’s Law to hold for at least another two decades.¹

The problem with concentrated transistors is the heat they produce. The power required, hence the amount of heat that needs to be dissipated, goes up as frequency increases, and down as the inverse square of voltage. On average this heat has been increasing 17 percent annually

Figure S.1: Heat Density Trends



Source: Uptime Institute; www.uptime.com.

¹ www.webopedia.com/TERM/M/Moores_Law.html.

(see Figures S.1 and 1.4, Recommendation 1.4). As components shrink, this growth is compounded. Heat is beginning to threaten chips' existence—indeed, some chips threaten to melt themselves. Some researchers and experts even expect that within a few years, the heat produced by some chips will rival the heat density found on the surface of a nuclear reactor core.² Thus a major challenge for the high-tech sector becomes keeping data powered while keeping chips temperate.

Today, for every watt being consumed by a computer, roughly two to three additional watts are being drawn from the utility to cool the computer and provide it with protected power. In RMI's opinion, this formula is unsustainable: as computers become hotter and hotter, HVAC engineers will be forever "chasing their tails," as one engineer has described the situation, in an effort to keep CPUs from destroying themselves. In addition to wasting energy, increasing power density also increases the risks of system failure

and highly expensive downtime, and increases capital costs and construction delays.

As Dr. Wu-chun Feng noted: "The more power a CPU draws, the hotter it gets. The hotter a CPU gets, the more likely it will fail (or clock down), and likely, cause other components to fail. More concretely, Arrhenius' equation (when applied to microelectronics) predicts that the failure rate of a given system doubles with every 10° C (18° F) increase in temperature. And in fact, unpublished empirical data from two leading vendors indicates that the failure rate of a compute node does indeed double with every 10° C increase."³

Society can no longer afford "brute power" solutions. The maxim "if brute force isn't working, you aren't using enough of it" is outdated and inappropriate. If our species is to survive on the planet at anything like our current standard of living, we must learn to use energy and other resources as efficiently as nature does. This is a central goal of RMI's work.

It is important to note that while nature is incredibly efficient, it also routinely provides performance significantly in excess of comparable human systems. For example, human brains possess impressive capabilities, yet they are elegantly frugal.

Our ability to emulate *natural* energy systems is still very primitive. Energy consumption is usually a secondary or tertiary consideration in designing modern mechanical and electrical systems. The use of massive quantities of energy to force functionality is rarely, if ever, questioned. Greater performance at the cost of energy efficiency and system reliability does not make sense; there are better ways to achieve high performance. As Dr. Feng⁴ said, "While performance and price/performance are important metrics...the key metrics of this decade will be efficiency, reliability, and availability." These goals must, of course, be achieved within the context of reliably providing needed services and achieving needed goals.

² It is important to note that this number has not been verified by RMI.

The number has been passed through the computer community, and appears to have originated with Pat Gelsinger, Senior Vice President, Chief Technology Officer, Intel Corporation, who referred to the surface of a reactor in a presentation at the Intel Developer Forum, 28 February 2002 (see www.intel.com/pressroom/archive/speeches/gelsinger20020228.htm). Although the comment is viewed by some as figurative or garbled (reactor power densities are normally expressed per cm³, not cm²), the fact that many leading computer designers use it as a representative figure is illustrative of how accurate they believe it will prove to be.

³ "The Bladed Beowulf: A Cost-Effective Alternative to Traditional Beowulfs" by W. Feng_, M. Warren_, and E. Weigle (feng, msw, ehew_@lanl.gov), Advanced Computing Laboratory and the Theoretical Astrophysics Group, Los Alamos National Laboratory, Los Alamos, NM 87545, undated, p. 3.

⁴ Technical Leader, the RADIANT Group, Los Alamos National Laboratory. See Appendix A for a summary of his presentation.

⁵ **Charrette**: a very intensive, highly integrative, trans-disciplinary, roundtable workshop that brings together stakeholders and experts at

High-Performance Data Center Charrette⁴

On 2–5 February 2003, Rocky Mountain Institute (RMI—see www.rmi.org and p. 85) convened an innovative whole-system design charrette⁵ to challenge the current model for data centers and develop a groundbreaking data center design—a clean-sheet model with no compromises. We brought together about ninety high-level, broad-thinking, technically deep industry experts with strategic visions in engineering, computer design, energy, data management, business, real estate, and facilities design to consider the challenge of data center energy consumption. Our goal was to leapfrog past incremental and individual efficiency improvements to advanced whole-systems design models—models that ultimately cost less to build, work better, and save astonishing amounts of electricity (or are even net generators of electricity).

By digging deeply into questions of technology choice, system integration, and business strategy, we discovered numerous and significant benefits for owners, developers, and designers of data centers, manufacturers of site infrastructure, computing equipment and components, utilities, and related industries.

The result of this charrette was a data center design concept that reduces energy demand by an order of magnitude (89 percent) compared to today's standard designs, while providing equivalent computing power and greater reliability.

The charrette addressed several important issues:

- **Unnecessarily high energy bills.**

By following the recommendations in this report, data centers can reduce their energy requirements by as much as 89 percent in the future. Using existing technology a 75 percent reduction of power consumption is feasible. Part of this potential can be captured in existing facilities.

- **Unnecessarily high capital cost.**

Applying whole-system design principles can increase energy efficiency while reducing capital cost—due to the complementary effects of integrated design and the correct sizing of systems.

- **Grid dependence.**

Whole-system design improves power quality and reliability, reduces dependence on the utility grid, and significantly increases overall efficiency. The ultra-reliable onsite power generation system recommended at the charrette relegates the utility grid to back-up status and makes possible the sale of net power and ancillary services *back* to the grid, while ensuring high reliability. This is a key competitive advantage.⁶

- **Utility distribution charges and delays.**

The systems recommended at this charrette can substantially reduce project lead times and completion risks. Rather than waiting for the utility to strengthen distribution capacity to serve new facilities or expansions, developers can offer the utility a compelling proposition—and generate net revenue by selling power and ancillary services back to the grid when and where they're most valuable. Utilities and data center operators' contracts often require certain availability of power on demand. Reducing this "holding capacity" will free resources, lower costs, and eliminate potential grid bottlenecks.

the very outset of a design or problem-solving process. It yields an ambitious design product, typically conceptual with some extension into early schematic design.

⁶ Review comment from Joel Swisher: "Utilities resist allowing both functions, due to protection and interconnection complexities, especially in networked (rather than radial) urban distribution systems." See Recommendation 4.4 for a discussion of this point, and www.smallisprofitable.org.

- **Risks for the owner/developer.**

Using the optimized designs recommended at this charrette, data centers can be cheaper and faster to build. Modular design allows construction of only the capacity currently required, while making future expansion simple and reliable. Lower operating costs and facilities boost the whole industry.

- **Community opposition.**

Breakthrough energy-efficient design solutions and benign onsite power generation can improve the environment, minimize community concerns, and expedite approvals.

- **Uncaptured opportunities for product sales** for equipment that contributes to implementing these integrated design solutions.



Charrette Process

RMI's charrettes are a design process developed by RMI and the American Institute of Architects (AIA) for the Greening of the White House. Charrettes apply a comprehensive and whole-system approach to design. This leads to integrated solutions with improved economic, energy, and environmental performance, simultaneously and without compromise. Charrettes have subsequently been used for numerous large-scale projects of many kinds, ranging from vehicles to process plants and from buildings to refugee camps.

This three-day integrated design charrette, orchestrated by RMI, focused on capturing energy and environmental performance improvements for each aspect of data center operations. The charrette emphasized facilitated dialogue in the areas of innovative and functional energy-saving design and engineering, and strategies to improve environmental results of design decisions. Innovative design of system components and integrated systems should significantly lower operating costs, and has the potential to reduce first costs as well.

RMI saw the need to host a forum on creating ultra-efficient data centers. The timing was right. A broad range of industry experts accepted the challenge and participated in the charrette bringing an enthusiasm and willingness to go beyond conventional thinking that far exceeded that of a single client-focused event.

In the conventional, linear design process, key people are often left out of the decision-making process, or brought in too late to make a full contribution. In this charrette, the participants brought an unusually wide variety of perspectives. Some were computer chip experts, others were HVAC specialists; we had server engineers and real estate specialists. This diversity contributed to the success of this trans-disciplinary re-think process. The fact that there was no client allowed participants to explore deeply the best possible scenario for each of the issues raised while minimizing proprietary concerns.

RMI's charrettes are usually client-driven—clients pay for an in-depth exploration of their issues. This charrette was unique in that it was not organized around a specific client, a specific location, or a specific process.

About This Report

This report summarizes the discussions that took place and actions recommended. It is organized to follow the compounding savings from the native loads back up-stream toward the power source. Participants narrowed their focus to six topic areas:

Part 1: Native Loads (CPUs, Servers, etc.)

Part 2: Computer Power Supplies

Part 3: Next Generation Cooling

Part 4: Cooling (Heat Removal)

Part 5: Facility Power Supply

Part 6: Operations

There are more than fifty major recommendations.

One goal of this report is to stimulate further examination of the various components of data centers and the energy they consume. It is essential that these components be designed and combined in an integrated—or whole-systems—fashion.

Whole-systems thinking is a process that actively considers interconnections between systems and seeks solutions that address multiple problems at the same time. Some refer to this process as the search for “solution multipliers” via a “vision across boundaries.”

The process of integrated planning and design, identifying the performance goals *up front*, is critical to achieving a good design. This allows the team to capture multiple benefits from single design features and to optimize overall data center performance.

It is important to recognize this unique integration process and *whole-systems* way of thinking when considering the use of the recommendations in this report. Many of them *cannot be considered in isolation* because their success and cost savings rely on the successful implementation of other recommendations.

The Internet has become an increasingly important factor in our economy. At this charrette we were able to take advantage of the current business slowdown to step back and critically examine current practices. We can expect aggressive growth of Internet-related facilities to resume in the future. When that happens, no doubt the ideas developed at this charrette and presented in this report will help to ensure orderly, profitable, and environmentally responsible growth.

How quickly will the Data Center of the Future be realized? We don’t know, but the early 21st century lull in the economy and the bursting of the late 1990s technology bubble have provided all who work with data centers, computers, and high-tech real estate a chance to do data centers right.

Time to Do Things Right

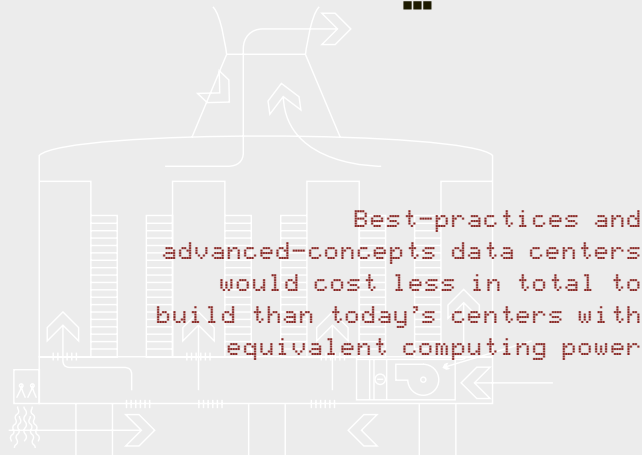
Question (Dale Sartor): Does the system designed by the Power Team use off-the-shelf technology, and if so, why is it not being used?

Answer (Brad Roberts): In the boom of building, I asked one of the building firms, ‘Why don’t you do this?’ They said, ‘We don’t have time to do that.’ In other words, they don’t have time to do it right.

Perhaps now, Roberts posited, we have the time.

We hope that readers will use this report as inspiration to challenge conventional designs of buildings, servers, CPUs, and support systems. But most importantly, we hope that you use it to challenge conventional thinking about energy consumption, and how we design and build systems around bits and bytes.

Summary



Rapid growth of “mission-critical” server-farm and fiber-optic-node data centers has presented developers and energy service providers with urgent issues. Resulting costs have broad financial and societal implications. Even in a sluggish economy where existing data centers (many, though, of low value) can be bought for pennies on the dollar, there are tremendous opportunities to significantly improve the performance of new and existing centers.

The RMI High-Performance Data Center Charrette produced design concepts that can reduce data center energy demand by an order of magnitude (89 percent) compared with today’s standard practice, while providing equivalent computing power, lower system capital cost, faster construction, and greater reliability. Using today’s existing technology, a 66 percent reduction of power demand is feasible. While this estimate applies primarily to new sites, many of the charrette concepts are also applicable to retrofits of existing facilities. While this estimate applies primarily to new sites, many of the charrette concepts are also applicable to retrofits of existing facilities. Figures S.2–4 quantify energy consumption when the recommendations contained in this report are implemented, and the matrix shown on pp. 26–31 classifies them by who needs to adopt them.

Figure S.2: Energy estimates by practice




Figure S.3: An estimation of HVAC, lighting, UPS, and computing power required in a redesigned 1-U computer
Absolute units (Watts) based on a single 1-U box for computing*

End Use	Current Practice	Best Practices with Current Products	Projected with Advanced Concepts	Assumptions
Computing (see Fig. S.6)	128	55	21	See Computing Worksheet
HVAC (see Fig. S.7)	103	24	5	See HVAC Worksheet; HVAC energy is computed as % of computing energy. But then where does HVAC efficiency show up?
Lighting (see Fig. S.8)	4	2	1	
UPS & Other (see Fig. S.9)	17	4	1	See Computing Worksheet
Total	252**	85	28	
% Energy Compared to Base Case	100%	34%	11%	

* To make all end-uses commensurable and normalize to W/s.f. (watts per square foot) units, this chart uses end-use percentages rather than W/s.f.
 ** Rumsey Engineers, Inc. (Oakland, CA), “Data Center Energy Benchmarking Case Study,” December 2002, Lawrence Berkeley National Laboratory. This study benchmarks the use of energy by data centers.

Malcolm Lewis¹ created Figures S.2–3 at the charrette to integrate the potential energy savings identified by the various working groups. The various components that make up these results are tabulated in the supporting charts on the following pages and are discussed in the summary text related to each group. The graph identifies scenarios for energy consumption reduction in existing data centers, data centers using best practices with current products (currently available technology combined in smarter ways), and a projection for advanced concepts not yet “on the shelf.” Each of these three scenarios provides the same computing power.

“Best Practices” and “Advanced Concepts” refer to the practices identified at this charrette and listed in this report. The first uses technology that exists today; the second uses new technology that must be invented or that exists but has not yet been put into common practice. The matrix shown in Figure S.3 assigns each recommendation to one of these two categories, and further categorizes them according to professional interests or disciplines.

One metric for comparing the efficiency of data centers proposed at the charrette is *total* utility

¹ President and Founder, Constructive Technologies Group. Dr. Lewis is a consulting engineer who specializes in design, energy analysis, and forensics of mechanical, electrical, and energy systems for buildings and industrial processes.

power delivered to the facility divided by the net power that goes directly into computing equipment. Using this metric for each scenario yields the results shown in Figure S.4.

The numbers shown in Figure S.3 show what Lewis called an obvious and expected “double-whammy” benefit of best practices and of advanced concepts. Because the energy required for data processing drops significantly as the efficiency of the computing devices themselves improves, the heat generated and the need to cool them decreases, often exponentially. Also important is pervasive oversizing, currently standard practice. It can cause the cooling-energy requirement to be as much as three times greater than is what *actually* required by empirical analysis. Thus, right-sizing of many kinds of equipment represents a huge opportunity. “Best practices” assumes that variable cooling infrastructure is in place—systems and controls that adjust equipment use according to a user’s needs, as well as that equipment’s supporting infrastructure (chillers, fans, etc.).

The capital cost (new and retrofit) of these efficient systems was not formally estimated; however, the cooling team calculated that an annual return on investment (ROI) of 35–400 percent is achievable through improvements to HVAC systems alone (see Figure 4.1.1: Cooling cost benefits, p. 51), and Lewis shares the view of many participants that both best-practices and advanced-concepts data centers would cost less in total to build than today’s centers with equivalent computing power.

Integrated planning and whole-systems design require that performance goals be identified *at the beginning of the process*. This allows the team to capture multiple benefits from individual expenditures and to optimize overall data center performance. It is important to recognize this unique integration process and way of thinking when considering the use of the recommendations in this report. Many of them *cannot be considered in isolation* because their success and their cost savings rely on the successful implementation of other recommendations.

Figure S.4		
Total Power / Computing Power =	With Current Improvements in Computing	Holding Computer Power Constant
Current Practice:	1.97	1.97
Best Practice with Current Products:	1.54	0.38
Projected with Advanced Concepts:	1.36	0.13

Whole-systems thinking actively considers the interconnections between systems and solutions that address multiple problems at the same time. Some refer to this process as the search for “solution multipliers.” It requires a “vision across boundaries” and “solving for pattern.”

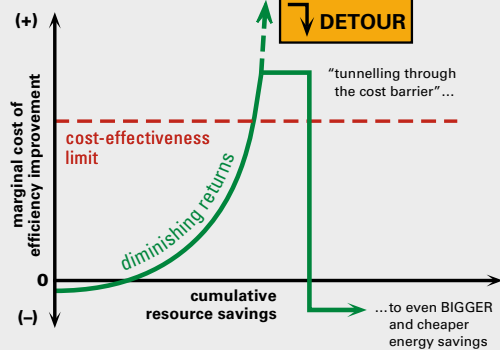
Comprehensively integrating topical recommendations is critical to achieving the best results. Design teams frequently use a value-engineering approach, instead of a whole-systems approach, for system or product selection. Value engineering is piecemeal; it prices design elements one by one to find the cheapest available product. But it fails to capture the benefits that can be achieved by recognizing that even though certain parts of the design may be more expensive, offsetting savings can make the *whole project* cost less and create greater value. Rocky Mountain Institute calls this phenomenon “tunneling through the cost barrier” (see Figure S.5). It is explained in detail in *Natural Capitalism*, especially in Chapter Six.

Reliability is the main critical element in data center facilities.

Therefore, efficiency cannot compromise reliability, and success will be facilitated if efficiency is shown to increase reliability.

The following paragraphs summarize the major findings of each working group at the charrette. This summary and the pages that follow describe how practicing energy efficiency can be achieved through combinations of advanced CPU and server technologies, smarter facilities, better HVAC system and power supply design, and more sensible metrics of what things really cost.

Figure S.5: Tunneling Through the Cost Barrier



Source: Rocky Mountain Institute

² Loosely quoted from *Natural Capitalism*, pp. 121–22.

Native Loads

To achieve the greatest possible energy savings in a data center, it is best to begin with native loads. This report is organized so that the reader can follow the compounding savings from these native loads back “up-stream” toward the electric power source. As *Natural Capitalism* states about reducing flow or friction in the pipes of a standard industrial pumping system:

Saving one unit of energy furthest downstream...avoids enough compounding losses...to save about ten units of fuel, cost, and pollution back at the power plant. Those compounding savings represent significant economic and environmental leverage...[enabling] each successive component, as you go back upstream, to become smaller, simpler, and cheaper. This...means that *downstream savings merit the greatest emphasis.*²

The first step in increasing efficiency is to recognize and account for the *full cost of each watt of power delivered* to the server. For data centers, paying average U.S. commercial electricity rates of \$0.07/kWh, this value is *at least* \$4/W—or twice that at least \$8/W, in places like Silicon Valley, New York city, etc., where electricity typically costs \$0.14/kWh. In particularly inefficient data centers, the value of each watt delivered to the servers can be as high as \$20/W. Note that power *always* saved (continuously) is generally worth several times as much as power saved intermittently, unless the intermittent savings come when power is especially expensive to deliver and this is reflected in its price.

Computing End-Use Energy

Today’s typical 1-U³ rack-mounted server uses approximately 128 watts. The conceptual “Hyper-server” developed at the charrette would be much smaller than current servers and would run on 21 watts and match the computing power of current practice 128-watt servers. Its electronic design is not radically different from that of the typical 1-U server; only its packaging differs. It achieves high levels of efficiency because designers have reexamined the equipment necessary for a server and have removed as much energy-intensive equipment as possible, notably superfluous fans and power supplies. Serendipitously, much current research is centered on creating low-power processors.

Dynamic resource allocation can solve the problem of unused system capacity by adjusting resources up and down as demand dictates, based on the true costs of those services. This can involve sharing resources across computers and software systems, across organizations, and even across the globe.

Figure S.6					
End Use	Current Practice	Best Practices with Current Products Assumptions		Projected with Advanced Products Assumptions	
CPU	70	20	Mobile CPUs	6	VLIW, low power, optimized CPU workload
HDD	10	10	Optimized for energy	3	Remote high-efficiency HDD
NIC/Misc.	10	10	Assume no change	5	Optimized for energy
Power Supply	33	10	Optimized for energy	6	One 3W rack-based power supply, high-efficiency; on-board converter (3W)
Fan	5	5		1	Fluid cooling
Total	128	55		21	×2 to achieve comparable performance, ×0.5 for savings from efficient resource allocation

Computer Power Supplies

One of the most important missing ingredients in efficient data centers today is efficient small power supplies. As much as half of the energy that enters a computer is wasted in its power supply. Both supplying and removing this wasted energy require significant amounts of energy and capital.

While there is no shortage of ways to cost-effectively increase their efficiency, power supplies are being optimized—to the extent that they *are* being optimized—using the wrong numbers. Most power estimates used today as the basis of design are severalfold wrong when actual costs

are accurately considered. Current thinking does not distinguish between component and system cost, nor between first and lifecycle cost. At a present value of at least \$4–8/W for each additional watt of server power, the data center is paying dearly for the inefficiency of the power supplies used in typical servers. If server purchasers were charged directly for the power and cooling loads they create, they would demand more efficient units from manufacturers. If designers and manufacturers understood that every watt saved is worth dollars per watt instead of cents per watt, they would build significantly more efficient devices.

³ 1.75 inches, the height of a pizza box.



Next Generation Cooling

Water can conduct 3,467 times as much heat as the same volume of air and requires an order of magnitude less energy to move a given volume. As temperatures on chips continue to rise and equipment loads continue to increase in density, liquid cooling becomes increasingly necessary. The first companies to move to liquid cooling will realize huge cooling-energy savings.

Resistance to liquid cooling may be greatly reduced if the liquid is kept well away from the chips by using non-liquid techniques to move heat from electronic components to liquid located off the board or even outside the rack. Changing the thermal path from convective (air) to conductive (liquid) eliminates the need for fans and minimizes the number of heat transfer steps in the thermal path. Successful implementation of liquid cooling systems requires standardization of plug-and-play cooling connections, locations, and methods for heat removal.

One alternative to liquid cooling is to use more real estate to reduce the *compaction*, or spatial density of heat sources, without necessarily changing HVAC technologies in a radical way. This can also save money (see Operations, p. 21).

Cooling

A 50–100 % increase cooling efficiency, corresponding to a 30–50% reduction in cooling energy cost, can be achieved with a ~40–400 percent *annual* return on investment (ROI)—with no decrease in reliability. Onsite cogeneration can improve reliability and increase chiller efficiency (by using waste heat) for larger data centers.

Higher levels of efficiency are achieved by more elegant and lower-cost solutions, such as air-side or water-side economizers and dry cooling. These solutions rely on the cooling potential of outside air whenever possible, minimizing use of vapor-compression equipment. Other high-efficiency solutions include evaporative cooling in dry climates (where data centers typically need humidification) and thermal-based cooling systems that use waste heat from onsite cogeneration to drive the heat removal process.

Equipment failure rates are three times higher at the top of the rack than at the bottom because that's where the heat collects. It would therefore make sense to put the most heat-producing units near the top of mixed racks.

Instead of operating data centers in the historically mandated 55–75°F range, improving the management of airflow and using new technology make it possible to raise the supply air temperature—to as high as 70–90°F—while increasing reliability and cooling system efficiency.

In large, densely populated data centers, the return air may embody larger total cooling loads (sensible + latent) than the outside air. In these cases, using outside air economizers will lower peak and average cooling loads. Data centers located in cool- and dry-climate regions can use natural cooling—which is free—by employing various techniques much of the year, including direct, indirect, and direct-indirect evaporative cooling, radiative or groundwater cooling, and enthalpy-controlled economizers.

Typically, data center ventilation systems are designed, installed, and operated at a constant rate for 8,766 hours per year. As a result, these systems frequently introduce far more outside air—that has to be conditioned—than is required. Except for command centers, few people continuously occupy data center critical space. Evaluating and minimizing ventilation rates can return big dividends in efficiency.

Chilled water systems with a capacity greater than 200 tons should operate at a total (supply fans through cooling towers) of 0.62 kW per ton. Systems greater than 60 and less than 200 tons should operate at a total of 0.83 kW per ton. These levels of performance have been achieved in real-world facilities. However, the full inclusion and commitment of all members of the design, construction, and development team are required to realize them.

HVAC

Figure S.7						
End Use	Current Practice (as % Computing Energy)	Assumptions	Best Practices with Current Products	Assumptions	Projected with Advanced Products	Assumptions
Heat Transfer out of Rack	0	Included in computing data for now	0	Included in computing data for now	0	Included in computing data for now
Heat Transfer out of Room	0.23	Air-based CRACs, constant volume, (2 W/cfm)	0.11	Ducted from racks to plenum, VAV, auto-rebalancing (1 W/cfm)	0.11	Fluid cooling or heat pipe to central fluid cooling system (assume 50% efficiency improvement over air)
Heat Rejection	0.58	Air-cooled, DX, poor part-load performance (2 kW/ton)	0.32	Water-cooled, chilled water, high delta T, optimized part-load performance, water-side economizer (0.62 kW/ton)	0.15	Fluid cooling or heat pipe to central fluid cooling system (assume 50% efficiency improvement over air)
Utilization of Waste Heat	0	None		BHCP with absorption cooling		Use waste heat to drive absorption cooling, plus BHCP
Total	0.81		0.43		0.26	
Question: How to handle recursion? Best-practice computing will have less heat load <i>and</i> have higher-efficiency HVAC. This table does not fully capture such interactions, and thus underestimates potential savings.						

Cooling (continued)

Optimizing existing control systems can provide a 20 percent reduction in total HVAC energy use on a typical system using only near-term, no-cost/low-cost solutions. A 30 percent reduction in total energy use is possible using variable frequency drives (which has a capital cost) + low-cost/no-cost. One of the simplest ideas—yet a concept with multiple benefits—is to network

CRAC unit controls in order to optimize and economize cooling efforts, and to allow the CRAC units to cool selected zones independently of other areas.

In the future, self-correcting, truly fault-tolerant control algorithms with automated adjustments based on measured data could remove human

error and lack of human responses to data. Building automation systems (BAS) could monitor rack/ chip temperatures and return air temperatures to optimize operating conditions and energy use. And dynamic management tools could deliver cooling where the data-processing load is, and/or move the data processing load where the cooling is optimal.

Lighting Energy

Figure S.8			
End Use	Current Practice (as % Computing Energy) Assumptions	Best Practices with Current Products Assumptions	Projected with Advanced Products Assumptions
Lighting	4.0 Over-lit, uncontrolled, in lightly-loaded data center	1.0 Reduced lighting levels, occupancy sensor controls; modern lighting equipment; zone to illuminate only areas of data center being used; in fully-loaded data center	0.5 Assumed further improvements in lighting efficiency. Visual effectiveness can improve with lower lighting levels through better lighting design (less veiling reflection and discomfort glare)

Facility Power Supply

The facility electrical supply system is a critical part of data center design, as it drives capital cost, operating cost, and the essential criterion of system availability.

The standard industry measure of reliability—five to six “nines”—is an incomplete measure. In data centers, even short interruptions can result in long computer downtime, data loss, and significant revenue penalties. Thus the rate of failure or MTBF could be far more important than the power supply availability or duration of outages.

It is important to note that the results of this charrette indicate that a data center could operate from a utility or onsite-generator supply voltage of 600V AC or less.

The charrette’s Power Supply Team recommended an onsite AC power distribution system. The choice of AC over DC appears to be as much a cultural as a technical partiality, however, and the group analyzed both AC and DC options. A differently composed group with more telecommunications switch experience might have recommended a DC solution.

The primary power supply should be an onsite generation system with minimum double redundancy, using the grid as backup. The recommended design eliminates 50 percent of the losses of today’s systems. More efficient than the grid, this system uses its waste heat to power a thermal-based cooling system, further reducing overall electrical demand. The synergy between the data center’s requirement for reliable, onsite power and the ability of onsite generation to satisfy the data center’s tremendous cooling requirement

simultaneously is a key strategy for reducing overall power consumption.To add capacity as the size of the data center increases (modularly), single modules can be added as necessary.

At least at present, the recommended system should be connected to the grid to ensure reliability. Ideally, unused capacity could be sold back onto the grid to keep generators running at full load, thus making them optimally efficient and shortening the payback period of the total investment. Unfortunately, the combination of power export and high-reliability operation is problematic.



UPS & Other End-Uses

Figure S.9			
End Use	Current Practice (as % Computing Energy) Assumptions	Best Practices with Current Products Assumptions	Projected with Advanced Products Assumptions
UPS Conversion Losses	13	7 Reduce over-sizing inefficiency	5 Adopt different technology for conversion and storage

Facility Power Supply (continued)

An optimally cost-effective system requires the capture of *both* the reliability benefits of standby operation and the energy savings of parallel operation. Although technically possible, it is difficult under present conditions to design *both* for power export to the grid and for premium reliability by island-mode operation during grid outages. Most distribution utilities will strongly discourage such a configuration. Thus, it is more practical today to design for premium reliability by island-mode operation during grid outages, and for parallel operation under normal conditions without the capacity to export to the grid.



Operations

There are as many opportunities to improve the performance of data centers by correcting the perverse incentives governing space, power, and cost relationships as there are by improving equipment and systems. The way to capture these opportunities is to “make true performance and costs transparent, and get the incentives right.” Incentives must be powerful and relevant, education must be a part of all data center considerations, and disconnected sectors need to work in unison.

Agents all along the value chain need to measure and to pay for the costs of the resources that they demand. The current system of charging users only on the basis of square feet encourages higher density of use and hence energy consumption well beyond the optimum. Current real estate models (design and construction relationships, lease and incentives) generate perverse signals

because they do not reflect the true cost of the capital and operating expenses necessary to deliver electricity of the requisite reliability to the server. Aligning market incentives with *desired performance* should eliminate today’s perverse incentive structures. Instead of charging on a per-square-foot basis, data center developers, designers, and managers need to select from a diverse menu of interrelated incentives: per watt, per power density, per teraflop, etc.—whatever metrics are practical and efficient.

A major misconception in space-to-power-density ratios is that cost per unit of computation comes down as power density increases. If properly calculated, as briefly discussed above, the present-valued cost of supplying energy can be as high as \$20,000 per kilowatt. The major cost is in the infra-

(Operations *is continued on p. 24.*)

Editor's Guide to Translation

High-k: k =dielectric constant. "With the difficult integration of copper largely complete and the onerous effort to bring low- k dielectrics to the interconnect stack well under way, technologists are moving to what could be an even more challenging task: replacing silicon dioxide with high- k dielectrics as the gate insulator. High- k materials such as hafnium oxide and zirconium oxide exhibit a tendency to "trap" electrons. At the International Reliability Physics Symposium here last week, technologists engaged in a furious debate over whether mobility degradation and threshold voltage instability are problems intrinsic to all metallic high- k materials."

Source: www.eetimes.com/story/OEG20030408S0047.

Voltage islands: the voltage island concept can reduce power consumption substantially by allowing designers to build processors that vary their voltages across a chip. A single system-on-a-chip processor could be built to run one voltage in one area, such as a processor core, and a different voltage in the other chip components. It could also switch off areas that aren't in use.

Source: <http://news.zdnet.co.uk/story/0,,t269-s2111576,00.html>.

VLIW: very long instruction word. Describes an instruction-set philosophy in which the compiler packs a number of simple, noninterdependent operations into the same instruction word. When fetched from cache or memory into the processor, these words are easily broken up and the operations dispatched to independent execution units. VLIW can perhaps best be described as a software- or compiler-based superscalar technology.

Source: www.byte.com/art/9411/sec12/art2.htm.

ICE: internal combustion engine

SiGe: silicon germanium: There is a 4% difference in the lattice constants of Si and Ge, so if one is grown on the other, the layer is strained and must be grown below the critical thickness. This strain may be used to vary the bandgap energy, band discontinuities, and effective masses, split the valley degeneracy, and adjust numerous other properties of the material. SiGe material has substantially higher mobility than Si material. The major advantage of SiGe is that it is compatible with CMOS and hence devices may be designed which may be fabricated on a Si chip alongside CMOS and bipolar. Hence SiGe devices can have substantially faster performance than conventional Si transistors while still being produced on Si production lines. As the cost of production lines increases as line widths shrink, SiGe may be able to provide some solutions.

Source: www.sp.phy.cam.ac.uk/%7Edp109/SiGeBackground.html.

Si/ins: silicon/insulator. Some advanced integrated circuits are fabricated as silicon-on-insulator structures, which facilitate faster operating speeds, closer component spacing, lower power consumption, and so forth.

Source: <http://eecs.oregonstate.edu/~flf/6309950.html>.

Negafans: no fans.

MEMS: microelectromechanical systems, reproducing conventional moving-part functions at a microscopic scale, typically using silicon-photoetching techniques.

Source: <http://mems.colorado.edu/c1.gen.intro/mems.shtml>.

VFD: variable frequency drive, using variable-frequency inverters to adjust the speed of motors. This can be a major energy-saver in pumps and fans, because fluid flow varies with speed while energy consumption varies at the cube of speed.

3,467× heat cap/vol + 10¹× movement Δ :

on reasonable assumptions, water has 3,467 as much heat capacity per unit volume, and requires an order of magnitude less energy per unit volume to move it, than air does.

10¹: "about ten." Ten to the power one means "on the order of ten," *i.e.* (since it's logarithmic) from approximately 3 to approximately 30. Writing it this way avoids implying greater precision than is known.

PAX: see footnote #7 on p. 57.

α, ε: absorption and emissivity. Absorptivity measures how well a material absorbs solar energy. Emissivity is the ability of a surface to emit radiant energy compared to that of a black body at the same temperature and with the same area; high-ε surfaces radiate infrared rays better, hence dissipate more heat.

DARPA PACT: Defense Advanced Research Projects Agency. See Appendix T: "DARPA funds power-aware architecture development" by Stephan Ohr, *EE Times*, 17 August 2000.

NEBS: network equipment-building system requirements. Refers to a family of documents that apply to telecommunications equipment located in a central office. The "NEBS Criteria," originally developed by Bell Telephone Laboratories in the 1970s, are a set of specifications for network facilities equipment. Their purpose is to assure that equipment is easy to install, operates reliably, and efficiently occupies building space. The expectation is that physical configurations and compatibility of equipment with a set of environmental conditions will help to reduce product installation and maintenance costs.

Source: www.ul.com/nebs and www.arcelect.com/NEBS.htm.

Crustacean Eater's Guide to High-Performance Data Centers

by Amory Lovins

Editor's note: This jargon-heavy series of bullet points is part of a powerpoint slideshow delivered by Amory B. Lovins at the charrette. The intent of the visually delightful presentation was to outline entertainingly—but comprehensively—the correct approach for reducing the power required by a data center, using the consumption of a lobster as a metaphor. The presentation is reproduced with the benefit of a translation guide, and in the spirit in which it was offered at the charrette.

Eating the Atlantic lobster

- Big, obvious chunks of meat in the tail and the front claws
- A roughly equal quantity of tasty morsels hidden in crevices, requiring skill and persistence to recover
- Go for both
- Mmmm!

The tail: power consumption

- What's it worth to avoid a watt of power consumption and heat generation in a data center? ~\$10.3PV el + ~\$9.6–16.5 capital; say \$20–27/W⁴—more than for solar kWh
- Low-V (≤ 1 V), high- k , voltage islands, VLIW, SiGe, Cu, Si/ins; RLX now gets ~5–8× saving

⁴ See update to this discussion in Part 2.2, p. 43.

⁵ See Sidebar: “No Straightforward Answer” on p. 67.

⁶ See Appendix K: *Cool Chips Overview*, by Jim Magdych of Cool Chips, Inc.

- Dynamic power management like laptops
- Superefficient power supplies; DC bus?
 - Could greatly improve uptime and reduce heat
 - Aebischer *et al.*, Canton Genève, 11/02: 2001 data centers used about half as much electricity/m² for telco on a DC bus as for internet applications on an AC bus, perhaps partly due to that difference⁵

The front claws: heat transfer & cooling

- Innovative heat removal from devices
 - Negafans, VFD fans, MEMS fans, PAX laminar-vortex-flow fans (2× eff.) (see figure S.10), inkjets, micro-Stirlings, quantum-tunneling thermal diodes,⁶...
 - Diamond like film, carbon fiber, carbon nanotubes,...
- Water cooling? (could be dry-chip, plug-and-play; 3,467× heat cap/vol + 10¹× movement Δ)
- At least thoughtfully designed airflow!
- Extremely efficient air-handling and cooling
 - Passive, then semiactive, then active
 - Economizers, passive latent heat exchangers
- Heat-driven HVAC based on onsite power, system efficiency ~0.90–0.92, ultrareliable

Figure S.10: PAX fan



The morsels, scraps, and broth

- Building envelope, α , ϵ , shading, elevators
- Lighting (~1–3 W/m² when occupied, lights-off when just machines are present)
- What temperature and humidity range does the equipment actually require? (e.g., NEBS-compliant blades handle $>>50\pm 5\%$)
- Load diversity, thermal time constants. Hottest/most heat-tolerant units on top?
- Lots of little savings multiply: e.g., $0.9^{10} = 0.35$

The whole lobster: a fantasy

- Optimized system architecture/compilation: DARPA PACT aims at ~10–100× savings
- De-bloated code and pared overhead: more useful operations per instruction executed
- Optimally share/spread real-time workload, as with multiplexed chillers; why is total data-center load constant while work varies $\geq 3\times$? (e.g., Tadpole/Platform Computing)
- Comprehensive, radical device efficiency
- Superefficient heat transfer at each stage
- Onsite trigeneration (turbines, fuel cells,...)
 - Heat-driven HVAC; eliminate UPS and its losses
 - Just a simple gas-fired-ICE single-effect absorption chiller makes data center a net exporter of electricity

Operations (continued)

structure to supply the cooling and power. This leads to radically different conclusions about the economics of further technology compaction. This is mainly a cost of power *density*, so pricing per square foot and per watt can help spread the costs *and* power density optimally.

There are myriad disconnects between the narrow foci and missions of the individual sector specialists—real estate, facilities, finance, vendors, IT, and end users—and the best interests of the data center as a whole. All individuals involved in the planning, designing, siting, construction, operation, and maintenance of data centers need to share *goals* and *information* and any “*pain*” throughout all stages of the process. One sector should not be penalized so that other sectors might be rewarded; all should share in successes and failures related to energy consumption.

If people don’t know what something costs and do not have to pay for it, they cannot be expected to optimize it. Thus, it is important that we develop full and disaggregated cost assessments for equipment and electricity, and give them to agents/users/customers all along the supply chain. It is also important that we develop methods to calculate lifecycle cost/total cost of ownership. Using this information, private and public entities can make good decisions about computing, electricity, and other resources.

Performance-based design fees provide incentives that encourage design teams to create buildings and equipment that are optimally efficient by rewarding the team for the savings they generate *from the savings they generate*. Creating standards to measure efficiency provides incentives to improve efficiency.

Gathering and benchmarking operating data is another key recommendation of this charrette. Feedback on costs is essential both for operations (short run) and for planning (long run). Comprehensive and useful metrics must be developed and benchmarked. A list of recommended metrics was developed at the charrette.

Measurement and verification capabilities continue to improve rapidly while costs decline, allowing more cost-effective real-time monitoring and management of energy and buildings systems that can increase systems performance (including energy savings), improve system reliability, and increase mean time to failure.

Creating an independent organization to provide testing, experimentation, education, and demonstrations could produce significant improvements in cost-effective data center efficiencies. Many functions that such an organization could provide are discussed in this report. If necessary, it should be jump-started by state energy agencies that manage public-goods fees.

Conclusion

The charrette results clearly pointed out how quickly the value of saving one watt compounds throughout the entire data center. We detailed a reduction of 83.5 percent in the computing equipment itself. This translated into a 94 percent reduction in all the other building system loads that support the equipment loads. This amplification illustrates how the savings in one system cascades into numerous related systems, not only saving energy but also reducing equipment size, complexity, capital cost, and causes of downtime. Additionally, simply looking at energy consumption does not measure other operational costs, such as human costs and the lost revenue from downtime and unreliable performance—not to mention the simple cost of maintaining the systems. Finally, in the case of data centers, efficient design massively reduces the quantity of material resources needed to provide computing services.

Matrix of Recommendations

A three-day transdisciplinary workshop with ~90 industry experts synthesized ways to design, build, and operate data centers that would use approximately tenfold less energy, improve uptime, reduce capital cost, speed construction, and enhance the value proposition. Some of the >50 integrated recommendations can also be applied to existing data centers.

Matrix of Recommendations
Part 1

Recommendation	Available Now	In Development/Future	Equipment Designers/Manufacturers	Engineers/Architects	Building Owners/Developers	Real Estate Brokers	Policy/Industry Assoc. Trades	Information Technology People	Comments
Part 1 Native Loads (CPUs, Servers, Software)									
1.1 New design paradigm: “Hyperserver” offers significantly improved efficiency		X	X	X Specify				X	The “Hyperserver” concept developed at the charrette offers significantly improved efficiency. It would be much smaller than current servers and would run on at total of 21 watts for equivalent computing power
1.2 Define operating envelope	X		X	X			X	X	Nameplate data is essentially useless from a practical design standpoint. It grossly overstates HVAC load and the maximum peak electrical requirement of the electrical power supply, fully loaded. Manufacturers should report actual loads for both electrical and mechanical for a “standard” operating configuration
1.3 Reduce or eliminate heat sources and improve heat management	X	X	X	X Specify			X Push development	X Push development	Rethink what goes into a server and remove as much energy intensive equipment as possible, notably fans, disk drives, and power supplies. Develop alternative chip cooling strategies
1.4 High efficiency CPUs	X	X	X	X Specify	X		X	X	Continue progress on new chips that use software to make hardware run more efficiently
1.5 Remove disk drives from servers		X	X					X	Disk drives do not need to be on the “motherboard” or within the server. Operating systems can be kept in a computer’s RAM
1.6 Right-size optimally efficient power supplies.	X	X	X	X			X	X	
1.7 Remove power supplies from server		X	X	X			X	X	Either remove them from the computer room altogether or put on top of rack where their heat can be quickly exhausted
1.8 Dynamically allocate resources	X	X	X	X	X			X	There is currently a problem of unused system capacity. Improve use of hardware, OS, application, and storage systems by throttling resources up and down as demand dictates based on the true costs of those services. This can involve sharing resources
1.9 Create an Energy Star standard for servers		X	X		X		X Push development	X	Create an Energy Star standard requiring that servers default to a standby mode when not being used

Recommendation	Available Now	In Development/Future	Equipment Designers/Manufacturers	Engineers/Architects	Building Owners/Developers	Real Estate Brokers	Policy/Industry Assoc. Trades	Information Technology People	Comments
Part 2 Computer Power Supplies									
2.1 Recognize and account for the full cost of each watt of power delivered to the server	X		X	X Perf.-based fees	X	X	X	X	
2.2 Create clear connections between power supply design, system efficiency, and power cost, with incentives to support efficient solutions	X		X	X	X		X Align incentives	X	
2.3 Focus on finding continuous, not intermittent, power savings	X			X	X			X	Power always (continuously) saved is worth roughly three times as much as power saved intermittently
2.4 Establish industry standards to increase power supply efficiency		X	X		X		X	X	This could be a non-proprietary rating system administered by an independent national organization such as IEEE or national labs
2.5 Improve power supply design		X	X				X	X	
Part 3 Next Generation Cooling									
3.1 Create system standards		X	X	X			X	X	Successful implementation of liquid cooling systems—both near-term and long-term—requires standardization of plug and play cooling connection locations and methods for heat removal
3.2A Convective cooling operating in conjunction with liquid cooling		X	X	X				X	
3.2B Cool the entire rack rather than each server	X		X	X	X		X	X	
3.3 Hybrid approaches for near term	X		X	X			X	X	Scheme 1: heat pipe connects the processor to a liquid column cooling “bus”; the remainder of the server is air-cooled. Scheme 2: transforms an entire 1-U server’s fan based cooling into liquid cooling
3.4 Conductive thermal path to liquid for future		X	X	X	X		X	X	Changing the thermal path from convective (air) to conductive (liquid) would eliminate the need for fans and minimize the number of heat transfer steps in the thermal path—reducing cooling system

Recommendation	Available Now	In Development/Future	Equipment Designers/Manufacturers	Engineers/Architects	Building Owners/Developers	Real Estate Brokers	Policy/Industry Assoc. Trades	Information Technology People	Comments
Part 4 Cooling									<i>A 35–95 percent efficiency improvement can be achieved with a 40–400 percent/y ROI—with no decrease in reliability</i>
4.A Low Energy Ventilation & Cooling									
4.A.1 <i>Increase the temperature range of cooling air</i>	X		X	X				X	<i>Systems that boost air temperature differentials increase cooling system efficiency</i>
4.A.2 <i>Manage airflow to reduce energy required for cooling and ventilation</i>	X		X	X	X		X	X	
4.A.3 <i>Minimize air side static pressure</i>	X			X					<i>Reduce system resistances by making detailed improvements in dynamic flow paths and efficiencies</i>
4.A.4 <i>Maximize use of free cooling</i>	X			X					<i>Bringing in cold-climate outside air instead of cooling return air from inside the data center reduces cooling loads</i>
4.A.5 <i>Natural ventilation</i>	X			X	X	X	X		<i>Natural ventilation is suitable for data centers located in dry climates—either hot or cold</i>
4.A.6 <i>Demand-controlled ventilation</i>	X			X	X				
4.A.7 <i>Additional ideas</i>	X		X	X	X		X	X	<i>Look for opportunities to combine developments with other activities that can use excess heat; use high-efficiency fans; keep motors out of airstream; place air handlers on top of racks; convert CRAC units to VFD; duct CRACs to return plenum; balance supply flow to match load</i>
4.A.8 <i>Wish list for manufacturers</i>		X	X				X	X	<i>Variable speed fans on enclosed servers; more efficient fans on boards and in CRACs; more efficient CRAC units; better managed, dynamically balanced air paths within server boxes and racks; ability to run at higher temperatures; servers that have laptop-type power supplies on board each box</i>

Recommendation	Available Now	In Development/Future	Equipment Designers/Manufacturers	Engineers/Architects	Building Owners/Developers	Real Estate Brokers	Policy/Industry Assoc. Trades	Information Technology People	Comments
4.B Efficient Heat Rejection in Large Data Centers									
4.B.1 Tune humidification and dehumidification cycles on existing systems	X			X					
4.B.2 Evaporative condensers/cooling	X			X					
4.B.3 Design and install chilled water systems greater than 200 tons to operate at a total of 0.62 kW per ton	X			X					
4.B.4 Design and install chilled water systems greater than 60 and less than 200 tons to operate at a total of 0.83 kW per ton	X			X					
4.B.5 Microclimate specific recommendations for northern US and cold climates	X			X					
4.B.6 Use waste heat from on-site cogeneration to drive HVAC system	X			X					
4.B.7 Dessicant cooling	X			X					
4.B.8 Thermal storage	X			X					Thermal storage is recommended only when all other methods have failed to provide the desired and required load reductions (>10KSF)
4.B.9 Wish list for manufacturers		X		X					More efficient coils, fans; substitute polypropylene for PVC fill-in cooling towers; efficient counterflow cooling towers; more efficient pumps; controls that work; more accurate and stable humidity sensors

Recommendation	Available Now	In Development/Future	Equipment Designers/Manufacturers	Engineers/Architects	Building Owners/Developers	Real Estate Brokers	Policy/Industry Assoc. Trades	Information Technology People	Comments
4.C Control Strategies									
4.C.1 General low/no cost optimizations	X			X	X			X	Shut off reheat and modify humidity set-points; continuous commissioning generate maintenance alerts; raised-floor housekeeping/maintenance
4.C.2 Establish environmental standards for mechanical and electrical systems by room type, and control to least energy-intensive values.	X		X	X	X		X	X	
4.C.3 Low/no cost solutions: CRAC optimization	X			X					Connect or network CRAC unit controls
4.C.4 Low/no cost solutions: reconfigure controls on central air handlers	X			X					Duct static pressure control: optimal point(s); supply air reset; fully utilize economizer, where applicable; minimize ventilation during non-economizer conditions
4.C.5 Low/no cost solutions: reconfigure controls on central plants	X			X					Optimize chiller sequence; apply condenser water reset control schedule; cooling tower sequence; for conditions below 45°F wetbulb, fully utilize water side economizer; fully utilize variable-volume pumping
4.C.6 Mid-level to high cost solutions	X			X					Network CRACs; add low-energy humidification; for VFD-controlled CRAC's: match air-side output to load; etc
4.C.7 Future control systems		X	X	X	X		X	X	Self-correcting, truly fault-tolerant control algorithms with automated adjustments; dynamically manage cooling capacity to deliver cooling
Part 5 Facility Power Supply									
5.1 AC power distribution system	X		X	X	X				AC supply is traditional in data centers, DC in telecom switching centers. (But cultural rather than technical issue.)
5.2 On-site power generation	X			X	X				The primary power supply should be an on-site generation system with minimum double redundancy, with the grid as backup
5.3 Interconnect with utility	X			X	X				
5.4 Address barriers to self-generation		X	X	X	X				An optimally cost-effective system requires both the reliability benefits of standby operation and the energy savings of parallel operation. (Institute of Electrical and Electronic Engineers (IEEE) is working to develop a national standard for interconnecting distributed resources with electric power systems.)

Recommendation	Available Now	In Development/Future	Equipment Designers/Manufacturers	Engineers/Architects	Building Owners/Developers	Real Estate Brokers	Policy/Industry Assoc. Trades	Information Technology People	Comments
Part 6 Operations									
6.1 Intelligent resource allocation	X				X				Figure out what you really want to do, then find the cheapest, most direct way of doing that
6.2 Improve information available	X				X	X			
6.3 Align incentives with desired performance	X		X	X	X	X	X	X	
6.4 Benchmarking	X		X	X	X	X	X	X	Track efficiency levels. Feedback on costs is essential both for operations (short run) and planning (long run) of data flow and processing capacity
6.5 Write more efficient code		X	X		X		X	X	Eliminate “bloatware” and make code that allows chips to scale up and down
6.6 Submetering	X			X	X	X			Submetering end uses allows real-time feedback and adjustments to reflect real costs
6.7 Measurement and verification (M&V)	X		X	X	X				Allows more cost-effective real-time monitoring and management of energy and buildings systems to increase systems performance/reliability
6.8 Continuous commissioning	X			X	X				Implement and maintain a comprehensive “best practices” and continuous maintenance system
6.9 Create self-diagnosing/healing systems		X	X	X	X		X	X	
6.10 Virtual servers		X	X		X		X	X	
6.11 Optimization tools		X	X	X	X		X	X	
6.12 Miscellaneous		X	X	X	X			X	Apply energy DSM; increase modularity of all components; minimize administrative burden and transaction costs; create transparency
6.13 Education, outreach, and training	X	X		X	X	X	X	X	Create “Best Practices” manual based on existing technologies, case studies, etc.
6.14 Demonstrations		X	X	X	X	X	X	X	
6.15 Energy Star and LEED Ratings		X	X	X	X	X	X	X	Creating standards to measure efficiency will provide incentives to improve efficiency
6.16 Create an independent organization to provide testing, experimentation, education and demonstrations		X			X		X		

Recommended Design High-Performance

Efficient individual components on the Hyperserver, when combined, optimize the entire system's efficiency. Their designs are not radically different from the typical 1-U server; only their packaging differs. The overall system is much more efficient, however, as our new energy requirement of ~28W is doing the same work as our previous requirement of 128W. This ~100W reduction also reduces all the resulting loads. The compounded value of these savings is tremendous.



Part 1: Native Loads¹ (CPUs, Servers, Software, etc.)


To achieve the greatest possible energy savings in a data center, it is best to begin with an examination of the native loads. The following quotation from *Natural Capitalism* explains why. This report is organized to follow the compounding savings from these native loads “upstream” toward the power source. As *Natural Capitalism* states:

TO LEAP FORWARD, THINK BACKWARD: Much of the art of engineering for advanced resource efficiency involves harnessing helpful interactions between specific measures so that, like loaves and fishes, the savings keep on multiplying. The most basic way to do this is to “think backward,” from downstream to upstream in a system. A typical industrial pumping system, for example, contains so many compounding losses that about a hundred units of fossil fuel at a typical power station will deliver enough electricity to the controls and motor to deliver enough torque to the pump to deliver only ten units of flow out of the pipe—a loss factor of about ten-fold.

But turn those ten-to-one compounding losses around backward...and they generate a one-to-ten compounding saving. That is, saving one unit of energy furthest downstream (such as by reducing flow or friction in pipes) avoids enough compounding losses from power plant to end use to save about ten units of fuel, cost, and pollution back at the power plant.

Those compounding savings represent significant economic and environmental leverage... [enabling] each successive component, as you go back upstream, to become smaller, simpler, and cheaper. This in turn means that downstream savings merit the greatest emphasis. The reason is simple. In a chain of successive improvements, all the savings will multiply, so they appear all to have equal arithmetic importance. However, the economic importance of an energy-saving measure will depend on its position in the chain. Savings furthest downstream will have the greatest leverage in making the upstream equipment smaller, and this saves not just energy but also capital cost. Downstream savings should therefore be done first in order to save the most money.

Downstream-to-upstream thinking is thus a special case of a more general rule: Do the right things in the right order. For example, if you’re going to retrofit your lights and your air conditioner, do the lights first so you can make the air conditioner smaller. If you did the opposite, you’d pay for more cooling capacity than you’d need after the lighting retrofit, and you’d also make the air conditioner less efficient because it would either run at part-load or cycle on and off too much...Once you’ve done the right things in the right order, so as to maximize their favorable interactions, you’ll have very little energy use left: Successive steps will have nibbled away at it a piece at a time, with each improvement saving part of what’s left after the previous steps. The arithmetic of these multiplying terms is powerful.”²



Source: E SOURCE. *Drivepower Technology Atlas* (www.esource.com).

¹ **Native loads** are those loads that carry out the critical functioning of a system, as opposed to ancillary equipment that supports those loads by providing such services as cooling or power conditioning. In the case of a data center, native loads include the computers and telecommunications equipment that provide services to customers.

² From *Natural Capitalism*.

1.1 New design paradigm


Today a typical 1-U server (see Figure 1.1.5) uses approximately 128 watts “inside the box,” as shown in Figure 1.1.1. Server designers tend to work by themselves, selecting standardized (and inefficient) components from catalogues. This is a hurdle to changing server design.

The “Hyperserver” concept developed at the charrette offers significantly improved efficiency.³ It would be much smaller than current servers (as shown in Figure 1.1.2) and would run on 21 watts. The Hyperserver comes without its own drives; these are housed separately, in a more efficient-to-operate location (see further discussion in Recommendation 1.5: Remove disk drives from servers. Keeping the operating system local (using DRAM or Flash) was also recommended, but greater RAM energy is required to handle this new configuration and this change requires IT sector education.

Figure 1.1.1: Today's typical 1-U server	
CPU	= 70W (A Pentium Four-type)
Two disk drives	= 10W
3 fans	= 5W
Power supply	= 33W (74% efficient)
Misc.	= 10W
Total =128W	

To match the computing power of today’s standard 128-watt servers, two Hyperservers are needed, meaning a total power requirement of 42 watts. However, with an anticipated 30–50 percent further saving from using a resource allocation approach,⁴ the actual energy requirement brings the Hyperservers back down to 21–30 watts.

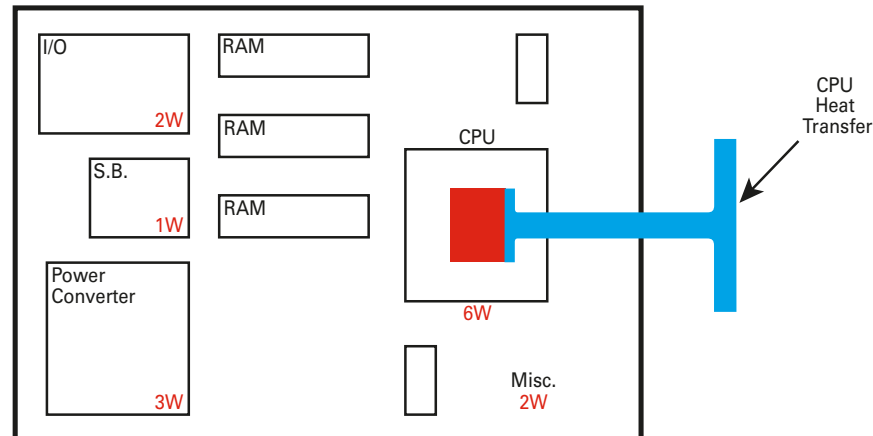
Figure 1.1.3: Hyperserver	
CPU	= 6W (See Recommendation 1.4)
I/O	= 2W
System Bus	= 1W
Power converter	= 3W (About 80% efficient)
Misc.	= 2W
Total = 14W on	
Plus offboard	
Disk drives	= 3W
Power conversion	= 2W
Liquid cooling	= 1W (Choice of this)
Total = 14W	
Grand total = 21W	



Of the 14W, the power converter(s) on the Hyperserver accounts for 3W. Although the fraction is high, in absolute terms the 3W are fairly easy to handle locally. The efficiency of conversion from AC to DC determines the heat generated in that process, based on the resistance of the rectifier and the voltage. If that efficiency is in the high 90 percent range, only a small amount of heat is generated as the current is quite low. This thermal energy can easily be managed by local air movement.

³ Daniel Lyons, “Super-Cheap Supercomputing?” Forbes.com, 25 March 2003, www.forbes.com/2003/03/25/cz_dl_0325star2.html. Star Bridge Systems claims to have created a reconfigurable “hypercomputer” that performs like a supercomputer but sits on a desktop, uses very little electricity, needs no special cooling systems, and costs as little as \$175,000. The secret is in the chips. Instead of yoking together hundreds or even thousands of microprocessors—as traditional supercomputers do—Star Bridge uses a dozen or so relatively inexpensive field-programmable gate array (FPGA) chips. Each FPGA can handle thousands of tasks at the same time, in parallel. Also, FPGAs can be reconfigured using memory cells connected to the transistors on the chip. So unlike most chips, an FPGA’s circuitry can be redrawn over and over again.

⁴ **Resource allocation:** “Basically we’re talking about powering down components when they’re not being used. At peak load, perhaps 20% of the processing power is being used in a data center. Resource allocation would power down 50% of the CPUs. It’s the same concept as CPU throttling, on a system wide scale.” (From the Heat Transfer Group’s discussion notes—see Appendix L for the complete notes.)

Figure 1.1.4: Hyperserver**In Data Center**

Disk-3W
Power Supply-3W
Fluid Cooling-1W

Server 14
In D.C. 7
2xTotal 42W

30-50% Savings
Power down with
resource allocation
21W

Figure 1.1.5: The Intel Itanium server

Source: Chris Hipp

Efficient individual components on the Hyperserver, when combined, optimize the entire system's efficiency. Their designs are not radically different from the typical 1-U server; only their packaging differs. The overall system is much more efficient, however, as our new energy requirement of ~28W is doing the same work as our previous requirement of 128W. This ~100W reduction also reduces all the resulting loads. The compounded value of these savings is tremendous. (See the discussion of the value of saving one watt in Part 2: Computer Power Supplies in this report.)

In his keynote presentation at the charrette, Dr. Wu-chun Feng remarked that even with an initially 50-75 percent higher cost per blade server, the 7-8× greater energy efficiency of the Green Destiny supercomputer yielded a ~90 percent saving on power, cooling, space, downtime, and system administration. This, in turn, yielded a ~65-75 percent lower total cost of ownership—exactly what this report is inviting data center developers to achieve.

In order to implement such a clever new solution, showcasing this new arrangement to IT people as well as consumers (education) is key.

In particular, it is important to show the sector and consumers that off-server disks are just as fast and reliable as CPUs with on-server disks, and that system security issues can be addressed. System security issues are easily handled today in the banking world. Here, the real issue that we are dealing with is the mindset that most data managers must have their own "unit."

1.2 Define operating envelope

“Nameplate”⁵ data are essentially useless from a practical design standpoint. They grossly overstate HVAC load and the maximum peak electrical ability of the electrical power supply, fully loaded (not realistic). Manufacturers should be encouraged to report actual loads for both electrical and mechanical systems for a “standard” operating configuration. This non-technical issue sounds simple, but is a very important step that should be accomplished before anything else—keeping in mind that there are a lot of other technical improvements noted in this report. There are several steps to the redefinition of the operating envelope:

- Determine the characteristics of operating conditions on the chip.
- Equip boards or racks with diagnostic tools that sense variations in temperature, humidity, and other critical conditions.
- Create better specifications for each component and system—including idling, median load, and maximum load—to replace exaggerated nameplate-based design and sizing errors.
- Define *reliability* as a function of operating characteristics and under varying loads. Reliability metrics include heat rejection, ambient temperature range, humidity, altitude, and voltage.
- Experiment with components’ configuration—they might be better left outside the server or even the data center. This could go a long way toward improving the way data centers are designed.

It is important to show the sector and consumers that off-server disks are just as fast and reliable as CPUs with on-server disks, and that system security issues can be addressed.

⁵ **Nameplate data:** The technical characteristics for equipment, as provided by the manufacturer. For computers and servers, nameplate energy data greatly exceed the power demand in typical actual operation.

Implementation

Foster development of an independent testing program (see Recommendations 6.15 –16.)

Fan Energy

A typical standalone low-end server normally has several fans. They are built into the power supply and also blow on the main processors. In the HP Netserver E800 model, for example, these fans have a total electrical power of 9.8 W. The power consumption corresponds to about 10% of the total power consumption of the server. In flat-built rack servers, the air for heat evacuation flows usually from the front to the rear. Highly loaded sections (processors) are often ventilated by air ducts. Several small fans (up to 10 or even more) are used for heat transport. For example, in a rack-optimized server (model IBM of xSeries 330) there are nine small fans with an electrical power of 2.5 W each. This results in a total electrical power of 22 W, which corresponds to about 25% of the power consumption of the server. The higher proportion in the rack-optimized server is due to its flat and compact construction with only small air ducts.


Source: *Energy- and Eco-Efficiency of Data Centres: A study commissioned by DIAE1 / ScanE2 of the Canton of Geneva*, by B. Aebischer *et al.*, 5 January 2003, p. 33. See Appendix S.

1.3 Reduce or eliminate heat sources and improve heat management

- Rethink what goes into a server and remove as much energy-intensive equipment as possible, notably fans and power supplies.
- Develop alternative chip cooling strategies (see Part 3 of this report).
- Operate external fans in conjunction with liquid cooling (see Part 3 of this report). This could eliminate the current need for numerous, less efficient, internal fans.
- Optimize heat sinks by using the appropriate fin orientation and design. (Charrette participant Chris Hipp told a story about having to rip out every other fin in a server rack because the fins were spaced incorrectly and were simply radiating and absorbing heat to/from each other, making heat rejection *less* effective.)
- Remove server box enclosures to increase airflow.
- Equipment failure rates are three times higher at top of rack than at bottom because that's where the heat collects. Most heat-intensive and -tolerant systems should be placed on top. At present, equipment is placed at random vertical locations.
- As much as 25 percent of the heat dissipated by a typical server is created by its numerous little fans (see Sidebar: Fan Energy, p. 36). RLX and similarly efficient servers have no fans. See further discussion of fan-less cooling options in Part 3: Next Generation Cooling.

Implementation

Research chips that can withstand higher temperatures and facilitate heat rejection. Note that the recommended improvements inside the server box make air temperatures tolerable. If these changes are adopted and the equipment loads are significantly reduced, there will be less incentive to switch to next-generation cooling, especially if a lower density real estate model is developed (see Part 6).



⁶ See Appendix A: Wu-chun Feng, M. Warren, and E. Weigle (feng_, msw_, and ehw_@lanl.gov), "The Bladed Beowulf: A Cost-Effective Alternative to Traditional Beowulfs" by W. Feng, Advanced Computing Laboratory and the Theoretical Astrophysics Group, Los Alamos National Laboratory, Los Alamos, NM 87545, undated p. 3.

⁷ See Appendix O: Transmeta press release: "Transmeta Announces Features of Next Generation TM8000 Processor for Energy Efficient Computing." (See <http://investor.transmeta.com/news/20030310-103475.cfm>.)

⁸ Floating operations (computations) per second.

1.4 High efficiency CPUs

"At idle, a Transmeta TM5600 CPU by itself generates less than a watt of power while a typical Pentium 4...generates as high as 75 watts. At load, the Transmeta TM5600 and Pentium 4 generate approximately 6 and 75 watts, respectively, while an Intel IA-64 generates over 130 watts! If the traditional mantra of "performance at any cost" continues, and hence, Moore's Law continues, the microprocessor of 2010 will have over *one billion* transistors and will dissipate over *one kilowatt* of thermal energy (see Figure 1.4)...."⁶

Much current research is centered on creating low-power processors.

- Continued development of chips that use software to make hardware run more efficiency (for example, Transmeta Crusoe and its successors)⁷ should be encouraged.
- Use laptop-style power management software.
- Reduce the power that the chip requires.

Implementation

Customers need to make it clear to chip manufacturers that power is a concern to them. Chip manufacturers need to produce convincing demonstrations to help customers get beyond the "faster-cycle-speed-is-better" mentality, showcase efficient chips to prove reliability, and create better metrics, including "FLOPS⁸/cycle," rather than clock speed.

1.5 Remove disk drives from servers

Disk drives do not need to be on the “motherboard,” or within the server. Operating systems can be kept in a computer’s RAM. These facts lead to important space- and energy-saving opportunities for servers that should be explored for data centers.

- Remove two drives per server and one power supply.
- Keep operating systems local (DRAM or Flash).
- Create a security system for shared resources. Shared larger drives are more efficient and more reliable. To do this successfully requires creating a security system that works with shared drives.

Implementation

Designers need to prove to the IT sector that servers with off-server disks are as fast, reliable, and secure as conventional servers. Performance standards for drives that specify watts per gigabyte might be one way to compare the devices.

Recommendation

1.6 Power supplies

- Right-size power supplies.
- Select optimally efficient models: select for the desired typical operating performance of the power supply integrated over its expected load range, not the full-load performance, and certainly not the nameplate load.
- Use two (or more) power supplies of different sizes to maximize efficiency for a given load.⁹
- Make supplies modular so they are hot-swappable; also, design them so they can be individually activated or de-energized, based on the total load, to optimize efficiency further.

⁹ Review comment from Tom Croda: “The systems view may optimize efficiency by using more small supplies that are distributed within a system. Each one runs a function and operates at maximum efficiency. Redundancy is gained by parallel processing with many processors. The loss of a power supply will result in the loss of a processor, but since there are many running in parallel, the loss of one will not cause a major failure.”



Photo courtesy Chris Hipp

1.7 Remove power supplies from servers

Remove power supplies from the servers, and either remove them from the computer room altogether, or put them on top of the rack(s) where the heat they produce can be promptly removed. The cost efficiency and modularity of the system may be better preserved by using a common DC bus voltage, with all conversions taking place on the board. Intuitively, one common supply voltage is preferable; then the board can step up to the two or more voltages needed.

The *multiple benefits* to be gained by removing the power supplies from the servers include:

- higher efficiency in a direct-bus approach;
- supply power at required capacity rather than overstated nameplate ratings;
- a more efficient power supply that can be custom-designed, improved, and optimized for lifecycle cost;
- removal of a major heat source from the board;
- cheaper equipment: buy fewer, far more efficient supplies;
- more reliable equipment: moves power, fans, and heat offboard;
- quieter: fans removed;
- size reductions, as removing components makes the board smaller; and
- reduces the need for redundancy by concentrating it; currently every server has redundant, low load (20–25 percent of nameplate), low-efficiency power supplies.

Recommendation

1.8 Dynamically allocate resources¹⁰

Currently, there is a problem of unused system capacity in most data centers. Dynamic resource allocation can improve the efficiency of hardware, the operating system (OS), applications, and storage systems by throttling resources up and down (in the short and long runs) as demand dictates, based on the true costs of those services. This can involve sharing resources across computers and software systems. Load shifting in the short term matches electricity needed with time of use, and in the long term optimizes computer resource acquisition and design. More specifically, dynamically allocating resources includes:

- Creating new software to optimize resource allocation across the board;
- Using demand forecasting to enable efficient management of processing and memory use on two levels:
 - Enterprise: easier to forecast demand, and
 - Internet: need “instant on” capability;
- Shifting computing load within the center or between centers;
- Scaling CPU power to data/load;
- Powering down CPUs (and other associated resources, including power supplies and fans) not in use;
- Sharing resources among users: CPU, disk drive, memory;
- Dynamically controlling HVAC and other end uses;
- Using variable speed, temperature-signaled server fan operation; and
- Optimizing servers and power supplies for estimated time-integrated loads.

¹⁰ Review comment from Geoff Wood: “Sun announced this very system within a couple weeks of the conclusion of the charrette, based on newly designed optimization software.” See Appendix N—Greg Papadopoulos, Chief Technology Officer, Sun Microsystems, “N1’s Computing-on-Demand to Drive Network Services.” *Network Computing Asia*, 1 February 2003. (See www.ncasia.com/ViewArt.cfm?Magid=3&Artid=18548&Catid=5&subcat=46.)

1.9 Create an Energy Star standard for servers

Create an Energy Star standard requiring that servers default to a standby mode when not being used. This alone would save a lot of money (and energy), particularly in smaller businesses where activity is rare late at night. This is difficult currently, however, because the only standby systems available today are for processors, not entire servers. Microsoft Windows currently supports standby mode only during lack of human activity. Intermittent traffic across a server could prevent sleep mode.

There is a way to solve this problem.

It is possible to design a front-end processor that awakens the server when it is specifically requested to operate, and to create various wake-up levels.

Note that hard drives take 2–10 seconds to awaken, but they are not the only mass storage devices with which this approach might be attempted. In some applications, solid-state memory may be competitive and would have almost zero latency time.

A low-power mode for servers means that power supplies need only provide very small amounts of electricity (a few percent or even less) compared to full-time, full-power operation. In this case the only way to get reasonable efficiencies is by providing a special DC power supply output to be used in this low-power mode, as is done in efficient power supplies for PCs.




Part 2: Computer Power Supplies

A central finding of the charrette is that the most important missing ingredients in data centers today are efficient small-scale power supplies (devices that transform electricity—delivered in bulk—into small amounts, useable by small machines). Power supplies are being optimized—to the extent that they *are* being optimized—using the wrong numbers. Most power estimates used today as the basis of design are severalfold wrong when actual costs are accurately considered. Current thinking does not distinguish between component and system cost, nor between first and lifecycle cost. If designers and manufacturers understood that every watt saved is worth (in present value) many dollars instead of cents, they would specify and build significantly more efficient devices.

Power supplies feed servers, fans, and routers. Servers generally use two power supplies, one of which provides backup in case the other fails. This parallel operation means that each power supply runs at a low load, decreasing its efficiency. Single power supplies in PCs typically run at 30–50 percent of their maximum design load, while coupled power supplies in servers run at only 15–30 percent load. This yields efficiencies around 50 percent, because power supplies are

designed to run most efficiently at well above 30 percent of full load (see Figure 2.1.1.). The output DC voltage of modern power supplies for servers is commonly 12 V. Further DC-DC transformation is needed to supply the required DC supply voltage for the equipment—*e.g.*, 1.5 V for modern processors—with further loss of efficiency.

The overall efficiency of the power supply system is, therefore, typically less than 50 percent. Fans are generally required to cool power supplies larger than 50 watts. This adds to their power demand, and increases the heat load to the space by the fanpower plus any net increase in power-supply losses.



Because server purchasers are seldom concerned about energy efficiency, cooling system costs, or site power supply costs, manufacturers do not design power supplies for efficiency

Despite the highly technological nature of a data center, power supplies for computers used in data centers are alarmingly inefficient. Most achieve, at best, a mediocre level efficiency, roughly 70–75 percent, at full load. However, as we have seen, they typically operate at only 15–40 percent of full load, a condition for which they are not designed or tested. At these load levels the efficiency is about 50–65 percent or lower (see Figure 2.1.1). In other words, as much as half of the energy that enters the computer is lost in the power supply. Supplying this lost energy requires a larger than necessary on-site power system, while removing this lost energy, which takes the form of heat, requires a larger-than-necessary HVAC system. Thus, both the power supply system and the HVAC system waste significant amounts of energy and capital.

To the extent that power supplies *are* designed for and tested at efficiency, they are typically designed and tested for maximum load rather than the part-load conditions under which most units operate. Power loss is considered at maximum load only so that designers may determine the size and cost of the heat sink. Ironically, that criterion is usually irrelevant because the average load is only around 20–25 percent of the maximum rating, so power supplies rarely operate at the full-load condition for which they are usually designed. Efficiency at actual, typical operating conditions is not considered and does not drive the design. Lowest initial component cost is the goal, not the lowest system operating cost or lifecycle cost. It appears that very few people even consider optimizing whole-system cost.¹ This is a major business opportunity.

There are many technical solutions for increasing power supply efficiency. Beyond a certain level of efficiency, these solutions can increase size and cost (see sidebar: “Size, Weight, and Cost”), but power supplies used in data centers rarely reach that point. Manufacturers design and build power supplies in large volumes to achieve the lowest initial cost. Because server purchasers are seldom concerned about or even aware of energy efficiency, cooling system costs, or site power supply costs, manufacturers do not design power supplies for efficiency, even when the cost premium would be minor.

While there is no shortage of ways to increase the efficiency of “power inside the server box,” there are some surprisingly important barriers to implementing efficiency measures. These barriers merit special attention.

For more discussions of the opportunities represented by power supplies see Appendix H: B. Aebischer and A. Huser, *Energy Efficiency of Computer Power Supplies*, and Appendix I: NRDC Report: *Power Supplies: A Hidden Opportunity for Energy Savings*.



¹ Whole-system cost is the integrated lifecycle cost associated with *all* of the components or elements in a system that affect each other; cost and performance. Costs are evaluated under expected operating conditions over the system's lifetime, and the evaluation considers their integrated functions. For example, high-performance components may cost more than standard components, but they may enable offsetting savings elsewhere in the system. Current thinking does not distinguish between component and system cost, nor between first and lifecycle cost.

2.1 Recognize and account for the full cost of each watt of power delivered to the server

The first step in breaking the vicious circle of efficiency minimization is to recognize and account for the *full cost of each watt of power delivered* to the server. For data centers this value is *at least* \$4/W undiscounted (average U.S. commercial electricity rates are \$0.07/kWh), while in places like Silicon Valley, New York city, etc., where electricity typically costs \$0.14/kWh, this value is *at least* \$8/W undiscounted.

These estimates are based on $1\text{W} \times 8766\text{ hours/year} \times \$0.07/\text{kWh} \times 0.001\text{ kW/W} \times 3\text{-year lifetime} \times 2\text{ W-input/W-delivered} = \$4/\text{W}$.

The last factor (2 W-input/W-delivered) is the “delivery factor,”² which is the ratio of total data center demand to the demand of the servers. We have used the conservative value of 2 for these calculations, but it can easily be more than 10. In data centers with a delivery factor of 10, due to compounding losses in efficient power-conditioning and HVAC systems, the value for each watt of power delivered to the servers is \$20/W or more undiscounted.

This calculation suggests that eliminating just 5 watts of loss on a 400-W power supply would save at least \$20–40. The power supply in a typical server costs about \$20–40 to manufacture. Thus, an improvement of less than 2 percent in power supply efficiency could save *more than the cost of the power supply*.

The savings would be even greater in a data center with a delivery factor higher than 2, which is not uncommon (but should be).

The costs calculated above are energy costs only. They do not include any system costs such as cooling capacity, power distribution, site power supplies, or maintenance. The additional capital cost of a highly reliable power system—additional UPS, batteries, switchgear, emergency generation, as well as the cooling equipment to remove each additional watt and the power supply for cooling—dwarfs the basic energy cost.³

Recommendation

2.2 Create clear connections between power supply design, system efficiency, and power cost, with incentives to support efficient solutions

Manufacturers are not concerned with efficiency because their customers rarely demand efficiency when they buy servers. Computer server purchasers are typically not charged directly for the power and cooling loads that the equipment they buy for a data center will incur. Their data center accommodations are usually priced per square foot occupied, regardless of the energy demand and cooling load they impose. Because server purchasers do not pay directly the capital and operating costs of the energy systems required to run and cool their equipment, they have no incentive to buy efficient units, and, as described above, server manufacturers have no incentive to make efficient units. Power supply manufacturers are even more remote from customer costs. If purchasers

had to account for the full cost of supplying power to servers, they would demand more efficient units.

At a minimum present value of at least \$4–8/W for each additional watt of server power demand,

(Continued on next page.)

² Delivery factor is an interesting and simple-to-understand metric: watts to the data center vs. watts to equipment. Members of the charrette’s computer Power Supply Group were somewhat skeptical that the cost could really be this high, so they spent a fair amount of time thinking about this number. They calculated it from basic numbers, and “reality checked” it with basic data center statistics. According to charrette participant Tom Croda, former Chief Power Engineer for Sprint, “Delivery factor ranges from about 1.5 at the best facility to 13 at the worst.” Thus the total cost of \$4–8/Watt for 3 years may be a significant underestimate.

³ To realize these savings fully (\$4–8/watt or more), energy needs to be saved all the time (~8,766 hours/year). See Recommendation 2.3 for more information on this point.

Recommendation 2.2 (continued): Create clear connections between power supply design, system efficiency, and power cost, with incentives to support efficient solutions

the data center is paying dearly for the inefficiency of power supplies used in typical servers. This cost is passed through to data center customers, so owners and operators do not see the need to care about inefficiency. Since server purchasers seldom consider the economic benefits of power supply efficiency, power supplies are being designed using the wrong criteria.



If server purchasers were charged directly for the power and cooling loads they create, they would demand more efficient units from the manufacturers. Once server manufacturers learn that every

watt saved is worth not pennies but many dollars, they will begin buying, integrating, and selling more efficient devices.

Size, Weight, and Cost

by Neil Rasmussen

A significant part of no-load losses can be eliminated without adding cost, weight, or size. They are a result of inattention to the design. Some circuits simply waste power for no reason. I have eliminated losses from existing production designs with no negative effect. When original designers are confronted, they invariably say, “I could have eliminated that loss but there was no reason to.”

But after the “free” loss reduction, cost and size increases will start to occur. Making copper traces⁴ and wires larger on printed circuit boards and coils takes up valuable space. Adding additional capacitors to reduce capacitor heating increases cost and size. Designing low-loss current and voltage snubber circuits, or circuits that recover energy that might otherwise be wasted, increases cost and size.

A detailed analysis could provide good cost data on light-load efficiency. One simple method for estimating the result, however, is to examine the variations between standard production power supplies. Assuming that the mean value from survey data is the average and that the best light-load efficiency supply is achieved with a reasonable design, then the best compared to the average is probably a good [minimum] estimate of what is available for free. If you are willing to spend \$1 per watt beyond that, in my engineering judgment, you could reduce the loss by 30% without difficulty. [The whole-system lifecycle value of saving a watt can be \$10–20/W. —Ed.]

⁴ See “Tradeoffs between Copper and Core Loss,” p. 46.

2.3 Focus on finding continuous, not just intermittent, power savings

Power always saved (continuously) is worth several times as much as power saved *intermittently*. Intermittent power savings occur when certain equipment is switched off. However, even when some equipment might be turned off, support equipment (including power distribution, HVAC, UPS, chillers, fans, etc.), required to cool it when running, is likely to keep operating undiminished. To the data center operator, these systems represent significant costs.

Eliminating a power requirement altogether (continuously) saves not only the cost of the power itself; it also saves the capital and operating costs of the support systems that are no longer necessary. This principle of whole-system or “integrated” design, developed by Rocky Mountain Institute many years ago, permits “tunneling through the cost barrier”⁵ so that very large savings often cost *less* than small or no savings.

Recommendation

2.4 Establish industry standards to increase power supply efficiency

Develop broad-based requirements and standards for industry declaration of power supply performance statistics, including efficiency-versus-load charts to show part-load performance. This could be an Energy Star standard, or it could be a non-proprietary rating system administered by an independent national organization such as IEEE or one of the national labs. Standards educate consumers about the performance of power supplies and help them appreciate the importance of no-load and part-load loss, not just full-load loss.

Recommendation

2.5 Improve power supply design

Under normal conditions (very low loads) the copper losses in a switching power supply are only a small fraction of the total loss. Increasing copper (wire) size has a more positive impact on the high-load efficiency of the power supply than on the low-load efficiency. Therefore, increasing copper size is not an important issue when prioritizing changes to improve efficiency design is. (See sidebar: “Tradeoffs between Copper and Core Loss.”)

Making improvements in the no-load losses would increase the efficiency 0–25 percent. Reductions of up to 50 percent in no-load loss can be achieved very cheaply while realizing a profit over the life of the power supply.

While size and weight are not always related to efficiency, there are some obvious situations in which increasing the efficiency of a certain part of the power supply will result in a weight increase (see sidebar: “Size, Weight, and Cost”). For example, most of the losses in the low power range are created during switching operations. Every time transistors in the power supply switch on or off,

a small amount of energy is expended in both the transistor and the transformer core. As the switching frequency of power supply increases, this small amount of energy is expended more times per second and results in larger and larger power losses. The loss in the transistor depends on the voltage and current present when it switches, how fast it turns on or off, and the characteristics of the device itself. The losses in the core are related to the physical size of the core and the material it is made of. Efficiency improvements can be linked to weight increases in some cases

(Continued on next page.)

⁵ See www.rmi.org/sitepages/pid116.php.

and not in others. However, to achieve maximum efficiency gains from all sources, data center developers should accept higher total-system weights.

One way to increase efficiency would be to lower the frequency. However, this requires a larger transformer core to handle the same amount of power. It would also require larger capacitors and other filter components. In some applications, a small size increase will not have a large impact; in others it will be infeasible.

Another way to increase efficiency is to design the supply so that the transistors only turn on or off when there is no current and/or voltage across them. This is called resonant switching, and it can achieve lower losses without increasing size. Higher-quality materials for the transformer core increase its efficiency at a given frequency without changing the size. Amorphous iron can virtually eliminate core losses.

These four approaches (larger transformer core, lower-loss core materials, lower frequency, and resonant switching) are largely independent and could be implemented together for maximum savings. The application will dictate how sensitive the design is to size increases.

Using small “point-of-use” DC-to-DC converters on each circuit pack and distributing redundant DC power to them (at, say, 54 VDC) will improve overall system efficiency. In this way, each card has a converter operating at maximum efficiency

and the main 54 VDC converters can operate in a range of 75–100 percent of rated load, offering both redundancy and efficiency. (See sidebar: “A Preference for DC Power Supply” on p. 69)

Tradeoffs between Copper and Core Loss

Tradeoffs between copper and core losses occur in linear and switching power supplies.

A transformer designed for lowest *component* (not system) capital cost will not offer the lowest loss.

An investment greater than the lowest *component* capital cost design can gain efficiency at full load, or at fractional loads, but typically not both. Copper loss drives full-load efficiency, while core loss drives light-load efficiency.

Optimizing on system (not component) capital *and* operating (not just capital) cost can yield a dramatically different outcome, however.

Because cost and heat at full load are the assumed design drivers, designers typically design the lowest-cost transformer that meets the full-load efficiency requirement. Such a design does not consider core loss unless it is a substantial fraction of the copper loss. Design algorithms in this case always attempt to move losses to the core. This results in a transformer with poor light-load performance. A very small change in the objectives of the optimization results in significant reductions in core loss for a small power-supply incremental cost. Whole-system cost to the customer may well go *down*.

Iron losses in transformers at tens or hundreds of Hz frequency can be reduced by more than an order of magnitude simply by using amorphous iron transformer laminations, as is commonly done in efficient power distribution transformers.

We have not calculated the quantitative watt savings possible per dollar spent, but transformer losses are a significant contributor to part-load inefficiency. This effect is present both in switching power supplies and in linear (so called “wall wart”) power supplies, but it is larger in linear supplies.

Part 3: Next Generation Cooling

Water can conduct about 3,500 times¹ as much heat as the same volume of air. While this obviously gives water a huge advantage in heat removal efficiency, use of liquid cooling is a hotly debated concept for facilities where critical electronic data are stored. As temperatures on the chips continue to rise and equipment loads continue to increase in density, however, liquid cooling becomes increasingly attractive, even inevitable (see sidebar: “Panel Comments on Cooling”). While the industry’s reluctance to move to liquid cooling remains very strong, the first companies to do so will realize huge relative cooling-energy savings and other whole-system savings—such as reduced equipment failures—resulting from cooling improvements.² The key issue is whether innovative design can overcome traditional concerns about leakage.

Resistance to liquid cooling may be greatly reduced if the liquid is kept well away from the chips by using various non-liquid-based methods to move heat from the electronics to the liquid in an off-board or even outside-the-rack location. Heat pipes, carbon fibers, and a few other non-liquid media can transfer heat to a backplane liquid system, eliminating the need to have the water near the chip. Separating the two helps to relieve anxieties about water and electricity mixing. At the charrette we learned that at least one major company—Intel—has seriously evaluated this concept³ (see sidebar: “Water Cooling Developments”). Intel’s design closely parallels the ideas discussed at this charrette.

A direct benefit of liquid cooling is the elimination of fans and their associated loads and costs. With design optimization it may be possible to have a primary liquid cooling system for the chipsets and power sources, with free airflow (evaporative cooling) operating in parallel for a synergistic effect. One alternative to liquid cooling is to use more real estate to reduce the compaction, or spatial density of heat sources, without necessarily changing HVAC technologies in a radical way (see further discussion in Part 6: Operations.)

Regardless of cooling system design, however, the first step remains reducing heat output. Improvements inside the server box and in the power supply recommended in Parts 1 and 2 of this report put less demand on air-cooling. Updating components could decrease heat output by as much as 40–45 percent (see discussion of “cool chips” in Appendix K). It may be that if these changes are adopted and the equipment loads are significantly reduced, there will be little incentive to switch to next-generation cooling.

¹ The heat capacity of water is 1 Btu/lbm-F° * 62.4 lbm/ft³ = 62.4 Btu/ft³-F°. The heat capacity of air is 0.24 Btu/lbm-F° * 0.075 lbm/ft³ = 0.018 Btu/ft³-F° at sea level. Water-to-air heat capacity ratio is 62.4/0.018 = 3,467 at sea level (4,000 in Denver).

² A suggested value was 60%.

³ See Will Berry, ME, and Stephen W. Montgomery, Ph.D., “Dockable Server Concepts,” Intel Labs, 25 February 2002. (Available online at: www.securitytechnet.com/resource/rsc-center/presentation/intel/spring2002/VSW188.pdf.)

Water Cooling Developments

Customers of collocation data centers tend to use compact rack servers in order to reduce costs for rented floor area. In these flat built rack servers, electricity to drive fans for heat evacuation becomes more important, *e.g.* 25% of the power consumption of the server. This higher percentage of power consumption for fans of the rack-optimized server is due to the flat and compact construction of the device with only small air ducts. Operators of managed data centers could think about using water for heat evacuation. Indeed, with increasing power density of processors, manufacturers of servers and racks envisage direct water-cooling. A leading rack producer has developed a bus system for cooling water, which feeds coolers of different processors (Reference: Wasserkühlung für Server. *Magazin für Computertechnik*, www.heise.de/ct, October 2002). With this direct water-cooling system it is possible to dissipate much more heat with less auxiliary transport energy than by air ventilation. But using water in an electronic system is critical in terms of security, and therefore will be avoided as long as possible.

Source: *Energy- and Eco-Efficiency of Data Centres: A study commissioned by DIAE1 / ScanE2 of the Canton of Geneva*, by B. Aebischer *et al.*, 5 January 2003, p. 36. See Appendix S.

Airless cooling (liquid) systems offer the potential to place racks back to back, and thereby utilize a single wall system for cooling. Alternating the hot and cold aisles thus becomes unnecessary; heat is removed via the liquid, not by airflow.

Panel Comments on Cooling

Dick Bourne: I just did a little calculation. If you're willing to accept that a seven-[°F]-degree temperature rise is reasonable for water going through one of these racks, then one gallon per minute will cool a kW. That means 11 GPM to cool an 11 kW rack. That's not a very large water flow, and it probably can be warmish water.

Ken Brill: There's no free lunch. If energy density continues going up, we're going to have to embrace some things. Liquid cooling is coming. If we're headed to 400 watts a square foot, [then] instantaneous, uninterrupted cooling systems will be essential because one temperature excursion can destroy a room's worth of equipment. That has not happened yet, but when it does and a customer has to write out a check for \$15 million for new computer equipment, this change will occur. We are in an industry that reacts to things.

David Schirmacher: I think you can get past reluctance to water. Obviously, IT people are inherently concerned about water in their data centers. The real problem is that there are ideas coming from all over the place. If you come up with a water-based or fluid-based solution that people believe is the future, and you can document the benefits of it, and include steps that mitigate risks, then maybe you can push this idea through. These solutions are very expensive, and they are usually short-lived solutions. Consistency is the limiting factor.

Ron Perkins: Maybe we ought to get out of this little box. When the cooling load exceeds fourteen watts a square foot, don't cool it, get rid of it. Remove most of the heat—three-quarters of it—with once-through air and throw it out the stack. After that, providing enough spot cooling to satisfy the rack plus a really efficient exhaust fan eliminate the need for all the little fans. This strategy requires a high ceiling space so that the hot air can move away from the machine, stratify, and be easily removed. This would be much cheaper than trying to cool 400 watts of thermal energy per square foot.

3.1 Create system standards

The successful application of liquid cooling systems—both near-term and long-term—requires standardization of plug-and-play cooling connections, locations, and methods for heat removal that allow manufacturers to standardize interfaces between servers and liquid thermal buses. There are multiple ways to design liquid-cooling systems and to connect servers to cooling systems.

If different manufacturers create their own proprietary connections, it will be difficult—if not impossible—to install a mix of servers into any given rack. Setting standards for all combinations of conditions is required before wide adoption of liquid cooling can be practical.

Recommendation


3.2 General

- Combine convective cooling with liquid cooling. For example, this can be done by using large, external fans instead of many inefficient internal fans.
- Cool the entire rack, not individual servers. For instance, a small plenum could be designed into the back of the rack by covering it with a single panel. A single large, efficient fan in the panel would ventilate all of the equipment on that rack, replacing the multitude of small, inefficient fans in each piece of equipment (see further discussion in Recommendation 4.A.2).

3.3 Hybrid approaches for near term

- In the first scheme shown in Figure 3.3.1, a heat pipe connects the processor to a liquid column cooling “bus”; the remainder of the server is air-cooled. Quick-connect couplings at the top or base of the column allow fast, efficient heat removal. Possible energy savings for this initial step are 10 percent for the removal of the processor fans and an additional 12 percent for the cooling of the whole unit.


Figure 3.3.1: Near-term solution
Hybrid: Cools processor with fluid remainder air-cooled



- The second near-term design (*a.k.a.* “Near-term airless,” Fig. 3.3.2) transforms an entire 1-U server’s fan-based cooling system into a liquid cooling system. This would be done using a horizontal heat transfer plate that is enlarged around the processor. This plate would attach to the liquid column cooling “bus.” Possible energy savings generated by this measure are 20 percent for the removal of all fans and 25 percent for energy savings in the cooling system.

Figure 3.3.2: Near-term airless

- Works with pizza boxes or blades
- Horizontal transition accessory for blades
- As in hybrid—fluid column connects at top or bottom



3.4 Conductive thermal path to liquid for future

Changing the thermal path from convective (using air) to conductive (using liquid) would eliminate the need for fans and minimize the number of heat transfer steps in the thermal path. The elimination of fans could reduce cooling system power consumption by 20–30 percent. Removing heat to a liquid loop might increase the efficiency of the cooling system by 15–20 percent. The physical size of the servers should also be reduced because the fans have been removed and the heat sinks are smaller.

The recommended long-term solution, shown in Figure 3.4, is based on a “cool wall” located near or on the back wall of the rack. “Hot spots” (e.g., processors) would be located close to or against that wall. Whatever heat conveyance system is then used must carry heat only a short distance. If blade-type servers come to dominate data centers, each blade could be manufactured with an onboard heat transfer fin that plugs into the cool wall (“plug and play”). Airless cooling (liquid) systems offer the potential to place racks back to back, and thereby utilize a single wall system for cooling. Alternating the hot and cold aisles thus becomes unnecessary; heat is removed via the liquid, not by airflow. Layouts could become considerably more compact if desired.

- Provide radiative and conductive metal fins to transfer heat from processors and other hot components to the “plug-and-play” heat risers. These fins could be water-cooled to eliminate fans, so that servers radiate heat to the cooling fins and water removes the heat from the fins.
- Integrate cooling systems with server control.
- Completely eliminate forced air; consider a free convection strategy in which servers are placed in a vertical or sloping position on the racks to allow free convection as air is warmed and rises from contact with server component surfaces.
- Improve components.

Figure 3.4: Long-term blade cooling

Airless: conductive and/or radiant “powerless”:
modular/separate from racks




Figure 3.4 demonstrates possible methods of practically implementing this concept.

Design elements for a long-term recommended solution include:

- Redesign boxes and blades so that major heat sources, such as power supplies and processors, are located as close as possible to vertical cooling elements such as “cool walls” or liquid column cooling “buses.”
- If power supplies remain on servers, locate them near the plug-in base to shorten the heat transfer path. If power supplies are concentrated on the rack, rather than on individual servers, locate them at the “leaving side” of the liquid cooling loop, be it air or water.



Part 4: Cooling

A 50–100 percent increase in cooling efficiency, corresponding to a 30–50% reduction in cooling energy cost, can be achieved with a ~40–400 percent annual return on investment (ROI)—with no decrease in reliability (see Figure 4.1.1). Onsite cogeneration can improve reliability and increase chiller efficiency (by using waste heat) for larger data centers. Figure 4.1.1 compares the efficiency of different heat-removal systems and indicates the efficiency that each can achieve.

Higher levels of efficiency are achieved by more “elegant” and lower cost solutions, such as air-side or water-side economizers and dry cooling. These solutions rely on the cooling potential of ambient air whenever possible (as a result of differential temperatures) with minimum use of vapor compression equipment. Other high-efficiency solutions include evaporative cooling in dry climates (where data centers typically need humidification) and thermal-based cooling systems, such as absorption or desiccant cooling cycles, which use the waste heat from onsite co-generation to drive the heat removal process. These options can be far more efficient than the conventional ones shown (see 4b.9) but their characteristics will be very site- and design-specific.

Figure 4.1.1: Cooling cost benefits




Figure 4.1.2: Cooling cost savings from more efficient conventional cooling systems

For a normal 10,000-sq.-ft. facility, including all cooling equipment (not just the chiller) operated 8,766 h/y and paying \$0.09/kW/h, with a 300-ton cooling load (1 ton = 3.52 kW thermal = 12,000 btu/h)

Technology	\$/ton	First Cost	Investment Premium	kW/ton (COP)	Cooling System Efficiency	Annual Operating Cost for Cooling	Annual Cooling Op. Cost Savings	Savings in Cooling Op. Cost	ROI%
CRAC	1,600	\$480,000	\$0	1.4 (2.5)	0%	\$331,000	\$0	0%	0%
Water-cooled CRAC or Evaporative Condenser Chiller*	1,680	\$504,000	\$24,000	0.9 (3.9)	56%	\$213,000	\$118,000	36%	492%
Rotary Screw Chiller	2,400	\$720,000	\$240,000	0.83 (4.2)	69%	\$196,000	\$135,000	41%	56%
Centrifugal Chiller	2,800	\$840,000	\$360,000	0.68 (5.2)	106%	\$161,000	\$170,000	51%	47%

* A water-cooled CRAC (computer room air conditioner) includes a remote cooling tower and a cooling water supply, often groundwater. A compressor is typically located in a floor-mounted unit with an evaporator fan and cooling coil. An evaporative condenser chiller is usually located on a roof or outside and produces chilled water that ultimately flows through a fan coil unit in the computer room. What makes the evaporative condenser efficient is an elimination of one heat-exchange process. The hot gas from the compressor is condensed in a coil located in an adjacent cooling-tower-like enclosure, where recirculated water flows directly over the hot gas coils while a fan pulls ambient air through the enclosure, evaporating and cooling the recirculating water. In this way, heat exchange losses in the shell-and-tube condenser of the CRAC are eliminated.

Part 4a: Low-Energy Ventilation and Cooling

Recommendation

4a.1 Increase the temperature range of cooling air

Systems that boost air temperature differentials increase cooling system efficiency. Instead of operating data centers in the historically mandated 55–75°F-range, a 70–90°F range is reasonable using new technology. This significantly lowers the power draw for cooling and increases overall efficiency. In cool and/or dry climates, economizer¹ operating hours and effectiveness would improve greatly at these higher operating temperatures. Because data centers require less latent cooling than office space, economizer cooling should also be attractive in a wider range of climates and for more hours of the year than for offices. Higher air temperatures would increase economizer potential even further.

¹ An [air-side] economizer is simply an airflow control scheme that detects those times when it is more efficient to use ambient (outside) air directly to cool the space rather than using HVAC equipment. This requires an ambient air temperature or enthalpy sensor and control logic to decide if the ambient air is cool and/or dry enough to provide useful cooling, and then increasing the makeup ventilation air flow rate to deliver this cool air to the space.


Recommendation

4a.2 Manage airflow to reduce energy required for cooling and ventilation

Approximately 80 percent of all computer rooms are small installations that occupy ~200–1,000 sq. ft. Typically these small spaces are served by overhead HVAC systems. As we have previously learned, equipment failure rates are three times higher at the top of a rack than at the bottom because that's where the heat collects. Improving


the management of airflow in these rooms can reduce this concentration of heat and make it possible to raise the supply air temperature—possibly from the 55–75°F range to 70–90°F—thereby improving efficiency and increasing the opportunities for economizer operations. Improving airflow through the equipment racks

Figure 4a.2.1




PLAN

Overhead ducted supply & return
(return could be plenum)



will increase reliability. This recommendation discusses some ideas for accomplishing these goals, many of which are equally applicable to large and small data centers.

Figure 4.a.2.2: Flexible barriers



Alternating hot and cold aisles: Figures 4a.2.1–4 illustrate various methods for establishing alternate hot and cold aisles that provide the temperature differentials necessary to draw heat efficiently through the computer equipment to the return air plenum. This arrangement also avoids mixing hot and cold air before the air passes through the equipment.

Flexible barriers for airflow management:

Using flexible clear-plastic barriers, such as supermarket refrigeration covers, to seal the space between the tops of the racks and the ceiling or air return location can control airflow while allowing flexibility in accessing, operating, and maintaining the computer equipment below. Figure 4a.2.2 shows cool air being supplied through an under-floor plenum to and through the racks, into a separated, semi-sealed area for return to an overhead plenum. This displacement system does not require superchilled air, nor that air be accurately directed. This approach uses a baffle panel or bar-

rier above the top of the rack and at the ends of the cold aisles (see Figure 4a.2.1) to eliminate “short-circuiting” (mixing of hot with cold air). These changes should reduce fan energy requirements by 20–25 percent, and could also save 20 percent of chiller energy. With an upflow CRAC unit as shown in Figure 4a.2.1, combining pairs of racks with a permeable barrier allows hot air to be immediately exhausted to the plenum. Unfortunately, if the hot-cold aisle placement is reversed, as shown in Figure 4a.2.3 (with the cold aisles

(Continued on next page.)

Figure 4a.2.3: Reversed aisles





Figure 4.a.2.4



being the ducted aisles), the working (human) spaces would be hot—at temperatures up to ~90°F—so that error should be avoided.

APC offers a rack that draws air from the front to the back and collects it in ductwork. This is the ideal airflow design, as equipment manufacturers are embracing front-to-back airflow. For existing racks, a ducting array could be positioned to draw rack exhaust air (overhead) and minimize bypass air losses. For new rack systems, a simpler (and less costly) configuration—a collection plenum and discharge collar at the back of the rack—is recommended. A centralized exhaust air system, configured as negative-pressure VAV, could be connected to each rack collar. A metered flow of air would be extracted from each enclosure and

ducted directly to the HVAC return. This would greatly improve cooling psychometrics (avoiding unnecessary humidification and dehumidification), eliminate bypass air issues, eliminate the need for small fans in each rack, and leverage larger and more efficient fans.

Figure 4a.2.4 addresses the issue of data centers that lack front-to-back airflow potential, so the hot-cold aisle concept doesn't work. Here, the recommendation is to provide flexible barriers to create a "cabinet," or enclosure, around the racks. A gap at the base of the unit allows inflowing air to be controlled; overall, this system mimics a chimney, as it exhausts heat out the top.

Recommendation

4a.3 Minimize air-side static pressure

Reduce system resistances by making detailed improvements in dynamic flow paths and efficiencies (this can also apply in part to short-term efficiency improvements). Specific recommendations include using oversized ducting, sweet bends, larger coils and heat exchangers, bypass coils, low-face-velocity coils (<200 fpm or <1 m/s), high-efficiency fans, premium-efficiency motors mounted outside the airstream, localized main plant, short paths, and VFDs² on CRACs/AHUs. Note that it may be necessary to increase the dynamic pressure in order to limit the cross-flow of supply air into the return air stream.

² VFDs = variable frequency drives; CRACs = computer room air conditioners; AHUs = air handling units.

³ **Design Day:** The set of weather conditions within a 24-hr. period corresponding to the worst-case scenario in which an HVAC system is designed to operate.

⁴ For further information see http://davisenergy.com/_prod.htm.

4a.4 Maximize use of free cooling

Most data centers seem to be designed as self-sustaining boxes without regard to the climate in which they are located. In today's typical large, dense data centers, the return air may embody larger total cooling loads (sensible + latent) than outside air. In these cases, using outside air economizers with differential enthalpy controls will lower both peak and average cooling loads. Data centers located in cool and dry climates can use "free" cooling from air-side economizers much of the year. Bringing in cold-climate outside air instead of cooling return air from inside the data center reduces cooling loads, even on the design day.³ In relatively dry climates, put evaporative pads on condensers, and even use direct evaporative cooling as the next stage after partial economizer cooling.

A related concept is the "night spray"⁴ system, which comprises a piping grid and water spray nozzles installed on a roof surface. The system chills water to approximately 5–10 °F below the minimum night air temperature for use in cooling operations the next day. The water is cooled first by evaporation as it is sprayed, and then by radiation as it sits on the roof and radiates heat to the cool night air. In Sacramento, Calif., for instance, this system can provide 250 Btu/s.f. per night of cooling capacity, plus the advantage of washing a PV array or a "cool roof" to keep it clean. The roof also lasts indefinitely, and fire-insurance premiums may be reduced. Other approaches to free cooling, such as using an economizer on the water side, are discussed below and in the sidebar "Adiabatic Humidification."

Adiabatic Humidification

by Jeff Sloan, P.E.⁵

I am the design manager for McKinstry Co, the largest design-build mechanical contractor in the Pacific Northwest. We build computer facilities and operate them 24x7 for owners such as Microsoft.

Technologies exist to use adiabatic humidifiers to provide free cooling in cooler weather while efficiently providing desired humidification for computer rooms. This is not generally understood or put into practice. In our climate the annual air conditioning cost of a large server site can be trimmed 50% to 75% by direct use of outside air when the weather is favorable. The savings come from turning off the

refrigeration compressors at those times. This method should be considered whenever the outside air gets cooler than the air that can be removed from the room. If the air can be taken from the rack, it would have universal application.

Air conditioning engineers have lots of experience with cooling and with humidification, but very little experience doing both at the same time. Assuming humidification is required, adiabatic humidification is the answer to the problem of how to humidify dry winter air.

Most manufacturers of computer room cooling equipment only sell isothermal humidifiers that boil water like a tea kettle. Using steam to humidify cool outdoor

air is a net energy loss—one step forward, two steps back. Adiabatic humidifiers use the heat from the servers to evaporate water. McKinstry has lots of experience in applying them, but every job is custom and the product is not commonly available.

A "hyper-airconditioner" for server farms would use high-efficiency mechanical cooling in the summer, direct adiabatic humidifier (blowthrough, with downstream sensor to control economizer), and VAV fans. A "hyper-computer rack" would have VAV exhaust fans ducted for single-pass cooling with the airflow controlled by the servers themselves.

⁵ Unfortunately, Jeff was not a charrette participant. Contact information: jeffs@mckinstry.com, P.O. Box 24567, Seattle WA 98124-0567, (206) 832 8342, www.mckinstry.com.

4a.5 Natural ventilation

Natural ventilation is suitable for data centers located in dry climates—hot or cold. As shown in Figure 4a.5.1, raising the building off the ground allows “ground” cooling to occur. Hot and cold aisles (managed airflow), combined with tall solar- and/or wind-driven chimneys, create temperature differentials that, in turn, create natural drafts that pull cool air into the system at ground level and exhaust hot air through stacks at the top of the building. Evaporative spray, as in classical Persian draft towers, could provide supplemental cooling when natural temperature conditions aren’t sufficient to create a differential. In very humid or hot and/or hot-and-humid climates, additional mechanical cooling might be necessary.

Figure 4a.5.1: Natural ventilation

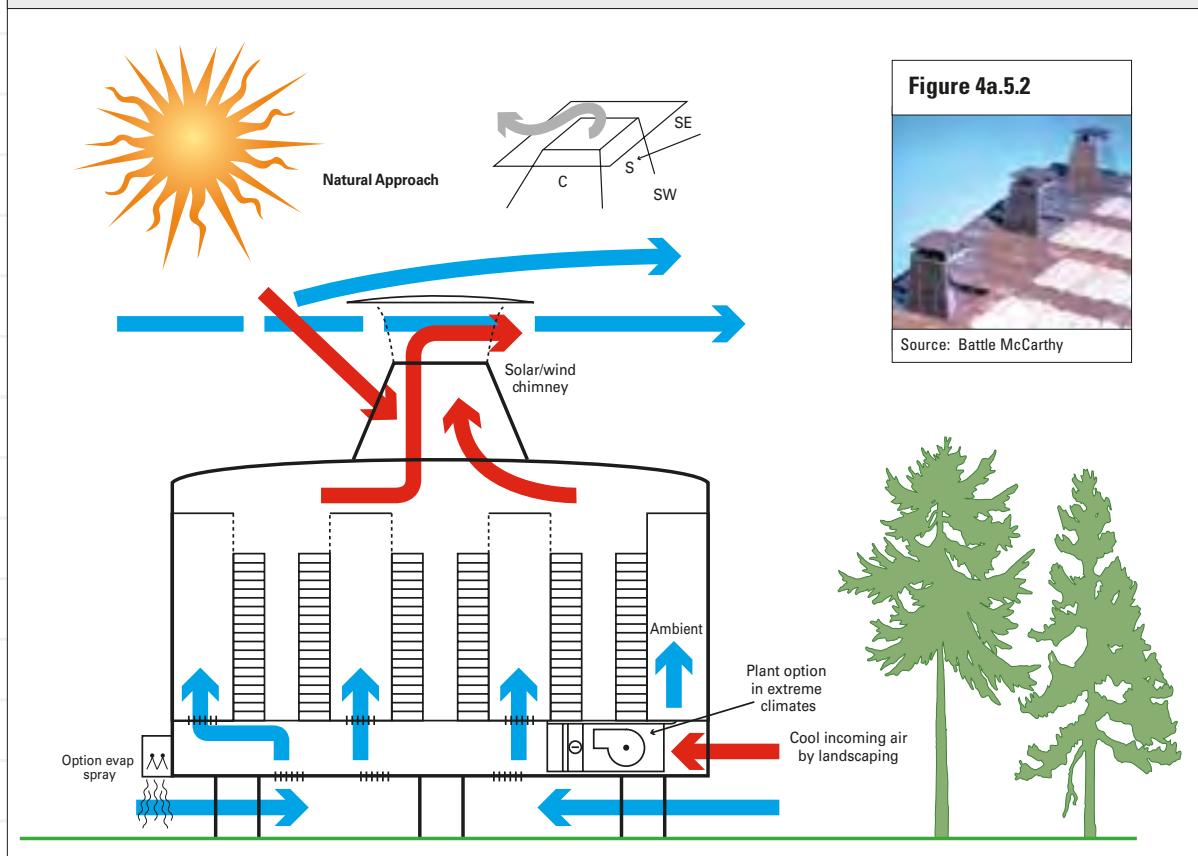


Figure 4a.5.2



Source: Battle McCarthy

Recommendation

4a.6 Demand-controlled ventilation

Typically, data center ventilation systems are designed, installed, and operated at a constant rate 24x7. As a result, these systems frequently introduce far-more-conditioned outside air than is required. Except for command centers, few people continuously occupy data center critical space.

Excessive ventilation imposes an additional cooling load on data center spaces during hot, humid weather, and can displace pre-conditioned air during the winter months (in colder climates) with drier air requiring additional humidification.

Therefore, evaluating and minimizing ventilation rates can produce big dividends in efficiency. Ventilation systems can run at very low rates [0.1–0.5 ACH⁶] and be triggered to increase when elevated carbon dioxide (more occupants/respiration) or VOC (from IT refresh or other contaminants) levels are detected. It is important to maintain a slight “positive pressurization,” however, to prevent air infiltration from adjacent areas.

⁶ ACH: air changes per hour.

4a.7 Additional ideas

- Look for opportunities to combine heating/cooling activities with other buildings or projects that can use waste heat generated by a data center. Examples include agriculture, laundry, restaurants, greenhouse, and swimming-pool heating and cooling activities.
- Use high efficiency fans, such as vaneaxial or mixed-flow fans. Watch for PAX impellers,⁷ a highly efficient new type of “propeller” that will be entering the commercial market in the next few years, initially ($\leq 200y$) in miniature computer fans.
- Keep motors out of airstreams, and buy the most efficient motor on the market (Motor-Master software from DOE shows no efficiency/price correlation up to at least 300 hp).
- Place air handlers on top of racks, next to the supply aisle to reduce ducting.
- Convert CRAC units to VFD.
- Duct CRACs to return plenum.
- Balance supply flow to match load.

⁷ For further information see www.paxscientific.com and *RMI Solutions* (www.rmi.org/sitepages/art7036.php). PAX Fan IT is rapidly commercializing the “biomimetic” PAX rotor for replacements on the small fans used in IT systems. RMI’s brief mention of this technology at the data center charrette may be well behind actual developments in the marketplace.

4a.8 Wish list for manufacturers

- Variable speed fans on enclosed servers.
- More efficient fans on boards and in CRACs (see 4a.7 above).
- More efficient CRAC units.
- Better managed, dynamically-balanced air paths within server boxes and racks.
- Ability to run at higher temperatures.
- Servers that have laptop-type power supplies in each box.



Photo courtesy Chris Hipp

Part 4b: Efficient Heat Rejection in Large Data Centers

Recommendation

4b.1 Tune humidification and dehumidification cycles on existing systems

- Remove the obvious inefficiencies:
 - eliminate bypass flow;
 - eliminate exhaust air entering equipment;
 - eliminate unnecessary latent cooling, since dry air is a greater problem for servers than humid air;
 - establish reasonable room temperature setpoints based on actual equipment and comfort requirements not mere habit;
 - ensure proper setpoints and calibration.
- Coordinate controls to ensure that one unit is not fighting with the other. Correcting this dismayingly common condition—easily detectable by measurement—can reduce power intensity from 2.4 (this may even be understated) to 1.4 kW per ton.



Recommendation

4b.2 Evaporative condensers/cooling

Retrofit existing air-cooled condensers and new chillers with evaporative coolers or cooling towers. There are no climate restrictions on this strategy; water always outperforms air, although the size of the savings will vary from cold climates to hot climates, and from arid to humid ones.

Recommendation

4b.3 Design and install chilled-water systems greater than 200 tons to operate at a total of 0.62 kW per ton

- Centrifugal chiller: 0.48 kW per ton, variable speed drive (*York or Trane*);
- Condenser water pump: 0.021 kW per ton (*3 gpm/ton, 30' TDH, 0.85 eff. pump, 0.92 eff. motor*);

⁸ Review comment from Amory Lovins: “Use MotorMaster (<http://mm3.energy.wsu.edu/mmplus/default.stm>) or Motor TechAtlas to select cost-effectively the most efficient motors possible. Do not oversize the motors. Use direct drive or high-performance belts (*e.g.*, Gates PolyChain GT with soft start or Habisat without).”

- Chilled water pump: 0.021 kW per ton (*2.4gpm/ton, 60' TDH, 0.85 eff. pump, 0.92 eff. motor*);
- Cooling tower: 0.011 kW per ton (*Ceramic, Tower Tech, or Shinwa cooling towers*);
- Air handling unit: 0.098 kW per ton (*400 cfm/ton, 1.5" TSP, 0.80 eff. fan, 0.90 eff. motor*);
- **Total = 0.62 kW per ton.⁸**

These levels of performance have been achieved on real-world facilities. However, the full commitment of all members of the design, construction, and development team is necessary to realize them. Careful negotiations with manufacturers are necessary to convince them to deliver equipment that performs at these levels, and to provide it at the lowest cost. Component and system performance must be specified, measured, and monitored using accurate sensors.

These values depend on achieving some specific objectives in terms of air and chilled water flow configuration, pipes and pumps, static pressure in each loop, etc. Some of the more critical assumed values are stated in parentheses after the target performance above. See E SOURCE (www.esource.com), *Space Cooling Technology Atlas*, for details.

Recommendation

4b.4 Design and install chilled water systems greater than 60 and less than 200 tons to operate at a total of 0.83 kW per ton

- Water cooled/water screw chillers: 0.68 kW/ton;
- Condenser water pump: 0.026 kW/ton;
- Chilled water pump: 0.021 kW/ton;
- Cooling tower: 0.012 kW/ton;
- Air handling unit: 0.086 kW/ton;
- **Total = 0.83 kW/ton.⁸**

Comments from Recommendation 4.B.3 above also apply here.

Recommendation

4b.5 Microclimate specific recommendations for northern US and cold climates

- Outside air economizer with heat recovery efficiency assumptions: 0.15 kW/ton (0.05 is the best noted in this design).
- At ambient temperatures below 55°F, it is possible to use a water-side economizer,⁹ possibly combined with thermal storage. Efficiency assumptions: chilled water pump 0.021; fan 0.1; cooling tower 0.02 = 0.14 kW/ton.

- At ambient temperatures below 40°F, it is possible to use a dry cooler-glycol/water system efficiency assumptions: 0.2 kW/ton. Reject heat to outside ambient by taking cold outside air to make 40°F glycol water; fan at the intake and at the discharge (used in cleanrooms of the northern United States).
- Earth-coupled vapor compression (ground source heat pump) efficiency assumptions: 0.02 for the pump, 0.9 for vapor compression, 0.1 for the fans = 1.02 kW/ton. A potential exists for eliminating the vapor compression (0.9) if 50–55°F (10–13°C) water is returned. If earth-cooled water (direct ground water) is taken directly to a coil, the result is 0.12 kW/ton.
- Eliminate the need for a chiller as a consequence of other measures, such as load reduction, outside air economizer use, ground water cooling, evaporative cooling, etc.
- Preheat outside air with recovered exhaust air (60°F).
- Use evaporative cooling/humidification (mist of de-ionized water).

⁹ A water-side economizer detects those times when ambient air is cool and/or dry enough to cool the chilled water directly using the cooling tower, rather than relying on vapor-compression equipment. The energy management control system decides if the ambient air is cool and/or dry enough to provide useful cooling. If it cannot, it then diverts chilled water to/from the cooling tower, increasing the flow when necessary, and bypassing the chiller compressor.

¹⁰ Review comment from Ron Perkins: “I would think that of the 450,000 Btu waste, we could recover about 80% of it, say 360,000 Btu/hour of it as 200°F hot water. This could drive an adsorption chiller to produce 21 tons of cooling at a COP of 0.7 (including cooling tower). You could boost the recovery efficiency to 90% and get a capacity of 24 tons cooling. Please note this the adsorption not absorption process. The more common absorption process gets a system efficiency of only 0.5 COP.”

Recommendation

4b.6 Use waste heat from onsite cogeneration to drive HVAC system

By making use of waste heat, onsite cogeneration can improve reliability and increase chiller efficiency for larger data centers. For example, a single, small (by data center requirements) 60-kW Capstone microturbine unit produces about 450,000 Btu/h, most but not all of which is recoverable. Converting that to cooling at a COP of a little less than 1.0 would provide about 37 tons of chilling.¹⁰

- An adsorber is a liquid-desiccant cooling system that uses waste heat to regenerate desiccant, and cooling towers to dissipate heat. It's like the absorption process, but better. Adsorbers have automatic load matching capability: chilling produced is directly proportional to waste heat generated.
- Absorbers produce steam and are less efficient than adsorbers in converting waste heat to cooling.
- If not using cogeneration, it is possible to shift the load through thermal storage devices, which can take advantage of off-peak utility rate structures (see next page).

4b.7 Dessicant cooling

Desiccant cooling dries the air, then humidity is added through evaporative cooling. It is applicable only to new designs, not retrofits. A disadvantage is that some excess heat is transferred into the incoming air. However the phase change of the water in the evaporative step can compensate for this. Efficiency assumptions: fan 78 percent; motor 94 percent = 0.24 kW/ton.

4b.8 Thermal storage

Thermal storage is recommended only when all other methods have failed to provide the desired and required load reductions. Its use applies predominantly to facilities larger than 10,000 sq. ft.

Water storage is preferred over ice because water is simpler, cheaper, and more reliable, although it requires more space. Use multiple tanks for system redundancy and emergency backup cooling potential.

Thermal storage could be linked to wet free cooling systems such as nighttime evaporative heat removal, bore holes, cooling towers, thermal ground or foundation coupling, and winter intake air path heat exchangers.

4b.9 Wish list for manufacturers

- More efficient coils and fans.
- Substitute polypropylene for PVC fill in cooling towers.
- Efficient counterflow cooling towers.
- More efficient pumps: the challenge is to get combinations of efficient pumps (84 percent) and motors (95 percent). The barrier isn't the technology; it's the failure to recognize and integrate the technology available.
- Controls that work: controls consistently fail to function as sensors drift out of calibration. Air conditioning manufacturers are starting to use better sensors. Presently the problem is that data center customers don't commission for efficiency / calibration (see Recommendation 6.8).
- More accurate and stable humidity sensors.

The above suggestions will yield the following percentage savings in cooling energy, using 1.4 kW/ton as the base case (air-cooled unitary CRACs):

- Water cooling the unitary equipment yields a 35% reduction in cooling energy demand;
- Chilled water system with water-cooled screw chillers between 60 and 200 tons=40% (*Centrifugal chillers are the best but are more expensive, screw chillers next-best and lower cost, reciprocating compressors last choice but still water-cooled.*);
- Chilled water systems with centrifugal chillers greater than 200 tons save 51%;
- Cogen. on the cooling side only saves 81% [*absorption chillers*];
- Desiccant cooling = 83%;
- Dry cooler glycol / water saves 85%;
- Cold climate: outside air economizer with heat recovery saves 89%;
- Water-side economizer of the chilled water system saves 90%;
- Earth-coupled direct water saves 91%.

Part 4c: Control Strategies

Improvements in control systems can provide a 20 percent reduction in total energy use on a typical unoptimized HVAC system using only near-term solutions (no-cost, low-cost) such as educating about and optimizing existing systems. A 30 percent reduction in total energy use is possible by adding VFD (which incurs a capital cost). For the most part, the recommendations listed next are near-term solutions using current technology.

Recommendation

4c.1 General low-/no-cost optimizations

- Shut off reheat and modify humidity set points (up to 3–5 percent savings on consumption at no cost).
- Continuous commissioning (track benchmarks); generate maintenance alerts when inefficiencies occur.
- Raised-floor housekeeping / maintenance: place tiles in right spot as required for proper airflow, eliminate unnecessary openings.

Recommendation

4c.2 Establish environmental standards for mechanical and electrical systems by room type, and control to least energy-intensive values

- Temperature;
- Humidity;
- Ventilation rate; and
- Lighting level.



Photo courtesy Chris Hipp

Recommendation

4c.3 Low-/no-cost solutions: CRAC optimization

One of the simplest ideas—yet a concept with multiple benefits—is to connect or network CRAC unit controls (*e.g.*, loop control) in order to optimize and economize cooling efforts, and give them the ability to work in zones. All CRAC units currently have the option of network connections. Networking would match the CRACs' collective output to demand and reduce “infighting” between units (*i.e.*, simultaneous heating and cooling in the same zone). Provide on-off setpoint control.¹¹

Recommendation

4c.4 Low-/no-cost solutions: reconfigure controls on central air handlers

- Duct static pressure control: optimal point(s);
- Supply air temperature reset;
- Fully utilize economizer cooling, where applicable; and
- Minimize ventilation during non-economizer conditions.

Recommendation

4c.5 Low-/no-cost solutions: reconfigure controls on central plants

- Optimize chiller sequence. This is usually adjusted manually in existing data centers. Load and necessary cooling should be matched.
- Apply condenser water reset control schedule.
- Cooling tower sequence (*i.e.*, one chiller, two cells).
- For conditions below 45°F wetbulb, fully utilize water-side economizer.
- Fully utilize variable-volume pumping.

Condenser water reset greatly reduces both consumption (kWh) and demand (kW) on chillers and CRAC units with compressors. Based on measurements from a reliable wetbulb transmitter sensing outdoor conditions, and following a tabular set of values, condenser water reset requires the tower (or fluid cooler or glycol cooler) water supply temperature setpoint to reduce (according to the reset schedule) as cooler, drier weather occurs. For example, when there is a wetbulb temperature of 50°F outside, the condenser water setpoint should be brought down to 65°F—which minimizes compressor energy—lowering energy consumption and peak demand in chilled-water or compression-cooled systems.

Recommendation

4c.6 Mid- to high-cost solutions

- Network CRACs (*i.e.*, loop control) to work in unison with on/off of unnecessary or redundant HVAC units.
- Add VFDs to all appropriate air-side and water-side devices throughout the HVAC system.
- Add low-energy humidification: replace electric steam generators with ultrasonic humidifiers, microdroplet spray, or other low-energy technologies. This is evaporative cooling with no reheat; it provides evaporative cooling on the way to reaching the humidity setpoint. Most importantly, this eliminates the parasitic heat gain of generating steam to humidify the space while providing “free cooling” in the process of reaching the humidity setpoint. (See also sidebar: Adiabatic Humidification, p. 55.)



¹¹ Group members were unsure whether or not this is typically done.

- For VFD-controlled CRACs, match air-side output to load.
- Install sensor/controls matrix, including temperature and relative humidity sensors, with sensors in-rack at inlets, to match load fully to HVAC output.
- For air-source DX¹² condensers: apply parallel VFD control of condenser fans.
- Install automated lighting controls: occupancy sensors (ultrasonic, not infrared), photocell-controlled dimmers, timers.
- Network controls to coordinate all these systems; fully utilize Ethernet and web-based platforms.

The energy-savings from VFD chillers are well-known. Coupling VFD chillers with reset-controlled condenser water (CTW) yields lower combined system power demand than can be anticipated from the manufacturer selection data. Caution: if CTW supply is too cold (*e.g.*, 55–60°F), the controls will speed up the VFD to compensate for low refrigerant head pressure, negating the “bottom-end” energy savings. Experience points to 65°F as the practical low setpoint.

¹² Direct expansion—packaged air-cooled air conditioning units.

4c.7 Future control systems

- Self-correcting, truly fault-tolerant control algorithms with automated adjustments based on measured data. Remove human error and lack of human responses to data.
- Provide networked IT temperature sensing at chip level as control point for HVAC.
- Take building automation systems (BAS) to the next level: monitor rack/chip temperatures and return air temperatures, and design the system to optimize operating conditions and energy use.
- Dynamically manage cooling capacity to deliver cooling “where the data-processing load is” and/or dynamically manage data processing to move the load “where the cooling is optimal.”



Photo courtesy Chris Hipp

Part 5: Facility Power Supply

The facility electrical supply system includes the UPS system, the backup generation units, and the switchgear and power delivery equipment. This equipment represents a significant share of the data center's total capital cost. It also is a major contributor to the electrical and HVAC energy losses, which represent a significant share of a data center's operating cost. Finally, the power supply system is one of the key failure modes, and it can place limitations on the availability of the data center. If power supply to *either* the computing equipment or the HVAC equipment is interrupted, a failure and downtime will occur. Thus, the power supply is a critical part of data center design, as it drives capital cost, operating cost, and the essential criterion of system availability.

Today, the industry standard is to maintain five to six “nines” (0.5–5 minutes of interruption per year—see Figure 5a) of reliability at the wall socket or its equivalent. Yet this is a very incomplete, even misleading, measure of reliability. In some industries, many short interruptions are tolerable but a single extended outage could be catastrophic. In data centers, however, the reverse is true. Even short interruptions, on the order of a few alternating-current cycles totalling much less than one second, can result in much longer computer downtime, data loss, and significant revenue penalties. Given the choice, a data center operator would far rather have one five-minute interruption per year than 300 one-second interruptions.¹



¹ For further discussion of reliability, see *Small Is Profitable* by Amory Lovins *et al.* (Rocky Mountain Institute, 2002), pp. 274–279, notably the last paragraph on p. 275; www.smallisprofitable.org.

There are more complete ways to measure availability and reliability.² Availability depends on both the frequency *and* the duration of failures:

Availability (A) = $1 - \text{MTTR} / \text{MTBF}$, where MTTR = mean time to repair (duration of outage) = $1 / r$ (r = rate of repair) and MTBF = mean time between failures = $1 / r + 1 / f$ (f = rate of failure).

Therefore, $A = 1 - f / (f + r) = r / (f + r)$, and Reliability (R) = $1 - \text{probability of failure} = \exp(-ft)$, assuming the failure rate f is constant.

For example, if $r = 1 / 12$ hours and $f = 1 / 20$ years, then R over one year is $\exp(-0.05) = 95\%$ (5% chance of a failure during one year); R over 10 years is $\exp(-0.5) = 60\%$ (40% chance of a failure during 10 years); and $A = 1 / 12 / \{1 / 12 + (1 / [20 * 8766])\} = 0.99993$ (4.5 “nines”).

In data centers, the duration of a power outage ($1 / r$) might be very short but still cause intolerable losses in terms of data and business.

In other words, the MTTR for the data center could be much longer than the MTTR of the power supply. This suggests that the rate of failure or MTBF could be far more important to data center performance than the power supply availability or the duration of outages.

Other metrics and indices are employed to characterize electric service reliability from a utility perspective, addressing frequency, duration, or extent of outages. The indices use *customer interruptions*, the sum of all customers experiencing all interruptions (customers with multiple events get counted each time) and *customer minutes* (the total of the product of customers interrupted times the duration of each interruption for all events). The most common indices in use are the system-wide average indices SAIDI, SAIFI, CAIDI, and CAIFI, all summarized in Figure 5b.³ The key for data centers is ensuring that even very brief interruptions are counted and minimized.

Figure 5a: Industry standard “nines”—a common but inappropriate measure of availability

Availability	Average Interruption Time per Year	Cumulative 24 Interrupted Hours Every
0.99 (2 nines)	3 days, 15 hours, 36 minutes	98 days
0.999 (3 nines)	8 hours, 46 minutes	2.7 years
0.9999 (4 nines)	53 minutes	27 years
0.99999 (5 nines)	5 minutes	274 years
0.999999 (6 nines)	32 seconds	2,740 years
0.9999999 (7 nines)	3 seconds	27,400 years

² See Allen, Whit, “Championing Power Quality and Reliability,” *Energy User News*, December 2002.
See www.energyusernews.com/CDA/ArticleInformation/features/BNP_Features_Item/0,2584,88122,00.html.

³ A thorough discussion of these indices can be found in Billinton & Allan (1996). A good discussion of their relationship (or lack of it) for DG planning is presented in Willis & Scott (2000).

In addition to reliability, the design recommendations in this section support the following design criteria for facility power supplies:

- scalability/modularity to allow easy additions to capacity without oversizing the initial design;
- doubled energy efficiency (reduce system and component losses by at least 50 percent); and

- competitive capital cost while reducing operating costs.

The model for these recommendations is a 10,000-sq.-ft. data center. It is important to note that the results of the “low power,” integrated, whole-system design championed at this charrette mean that the center could operate at an input voltage of 600V or less.

Figure 5b: System and customer metrics		
Metric	Computation	Description
SAIDI System Average Interruption Duration Index	$\frac{\text{sum of all customer interruption durations}}{\text{total customers in the system}}$	Average annual outage time per utility customer
SAIFI System Average Interruption Frequency Index	$\frac{\text{number of customer interruptions}}{\text{total customers in the system}}$	Average annual number of interruptions per utility customer
CAIDI Customer Average Interruption Duration Index	$\frac{\text{sum of all customer interruptions}}{\text{number of customer interruptions}}$	Average interruption duration. Sometimes referred to as average restoration time (ART)
CAIFI Customer Average Interruption Frequency Index	$\frac{\text{number of customer interruptions}}{\text{customers with at least one interruption}}$	Average number of interruptions for customers experiencing an interruption. All customer interruptions are included, but customers experiencing more than one interruption are counted only once
Availability	$\frac{\text{service available time}}{\text{total time}}$	Time fraction or probability that service is present
LOLP Loss of Load Probability	$\frac{\text{time that load exceeds capacity}}{\text{total time}}$	Possibility that load will exceed supply capacity
LOLE Loss of Load Expectation	expected time that load exceeds capacity, per year	LOLP in days per year time units, expected time that some interruptions will take place. Sometimes used instead to denote probability of experiencing at least one interruption per year
EUE Expected Unserved Energy	expected quantity of energy that would have been supplied during interruption	Sometimes referred to as “energy not supplied” (ENS) when used in historical reporting

5.1 AC power distribution system

After thoroughly exploring and evaluating numerous options, the charrette’s Power Supply Team eventually agreed to recommend an onsite AC (alternating-current) power distribution system—as opposed to DC (direct-current)—with short and simple distribution paths. AC supply is traditional in data centers, while DC is traditional in telecom switching centers. The Power Supply Team’s AC preference reflected its composition of nearly all AC experts; there was only one telecoms-oriented DC expert on the team. This appears to be as much a cultural as a technical issue. Despite the cultural preference, the group attempted to analyze both AC and DC options, but see sidebars: “No Straightforward Answer,” (p. 67) and “A Preference for DC Power Supply” (p. 69).

While DC is potentially more reliable, it is slightly less efficient when supplied from AC power,⁴ and is less widely understood than AC. DC is more practical at the personal-computer level, but less so at higher power levels. It is more difficult to provide DC circuit protection during faults or short-circuits. For example, the DC system’s associated with large UPS installations are very expensive and difficult to protect. DC can be more difficult to distribute, as the various DC-DC voltage changes require equipment that is not as familiar and inexpensive as conventional AC transformers

⁴ DC distribution is probably more efficient if its supply comes from a DC power source such as fuel cells.

and AC-DC power supplies. Some of the DC converters needed don't exist today, but could be developed. In the group's opinion, the larger wires and unconventional components would probably increase capital costs significantly. However, a whole-system analysis was not performed and might support a different conclusion, especially if time were allowed for the development of new components tailored to DC systems.⁵ In principle, an all-DC system could avoid many back-and-forth power conversions; its widespread use for telecoms switching centers cannot be assumed to be irrational; and its relative attractiveness may increase as efficient computing equipment decreases loads. Despite the group's conclusion, therefore, RMI considers the strategic choice of AC *vs.* DC distribution an open question ripe for further research.

See Recommendation 5.2 (on pp. 70–71) for a description of the recommended AC system. In the process of discussion, the team also diagrammed a possible DC system; this is shown in Figure 5.1.1 and discussed in the sidebars “Direct Current Facility” (at right) and “A Preference for DC Power Supply”⁶ (on p. 69). Its arguments appear to be un rebutted.

No Straightforward Answer

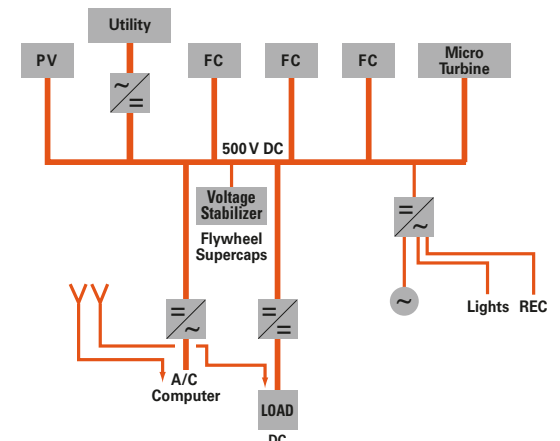
The typical electric load of a data center (equipment rooms only) is of the order of 400 W/m² for internet applications and 200 W/m² only for telco applications. Aside from different ICT equipment, another big difference is the way that the electric power is distributed among the individual ICT equipment. In a typical telco environment we find a central AC/DC transformation and DC-distribution, whereas in an internet environment most of the individual equipment have its own AC/DC transformer. This may be one of the reasons why telcos have a lower power load. Following Kolar (2002) there is no straightforward answer to the question of whether losses in power transformation could be substantially reduced by systematically using central AC/DC transformation. Further research is needed to evaluate the advantage of central AC/DC transformation and the disadvantage of higher transmission losses of DC-distribution.

Source: *Energy- and Eco-Efficiency of Data Centres: A study commissioned by DIAE1 / ScanE2 of the Canton of Geneva*, by B. Aebischer *et al.*, 5 January 2003, p. 36. See Appendix S.

Direct Current Facility

- 540V DC on main buses of a dual-bus system;
- extremely high reliability;
- ultracapacitors at the rack level provide secondary backup;
- more efficient, in the dissenting opinion of the sole DC expert in the group (see sidebar: “A Preference for DC Power Supply”);
- more costly (capital cost);
- having a DC power source increases efficiency globally and at the site (eliminates some inefficient equipment);
- mass market exists for DC components in the telecommunications industry.

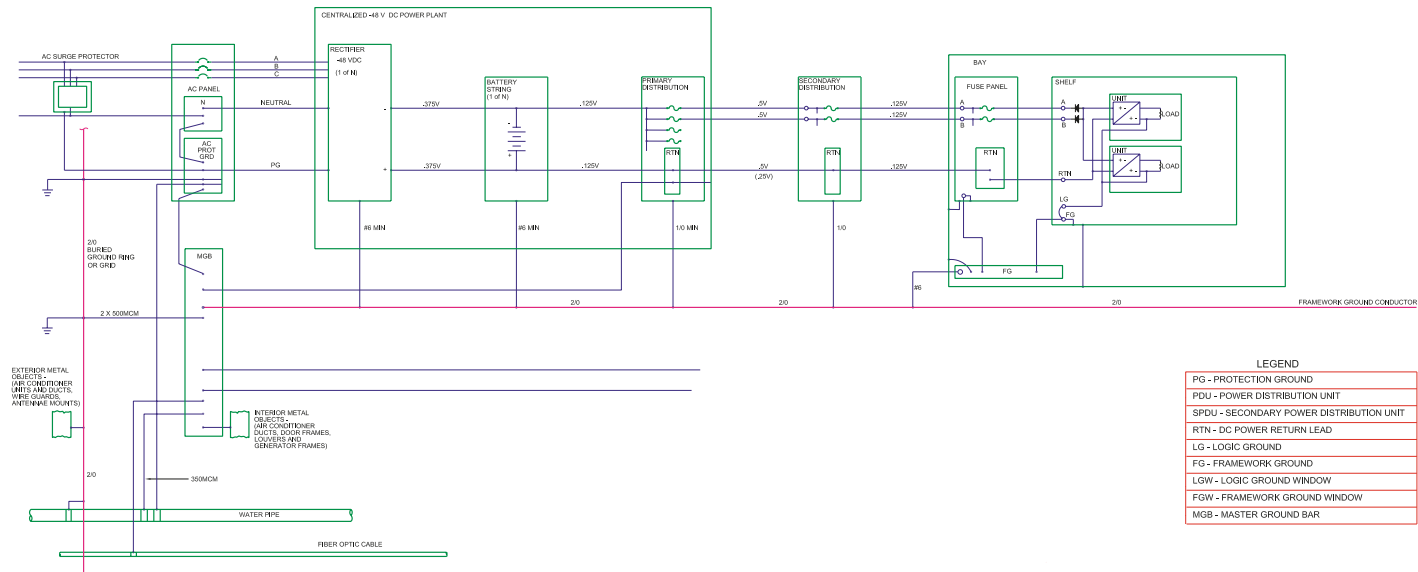
Figure 5.1.1: DC system of the future



⁵ Some of the inefficiencies of DC use in data centers are not known because no data exist.

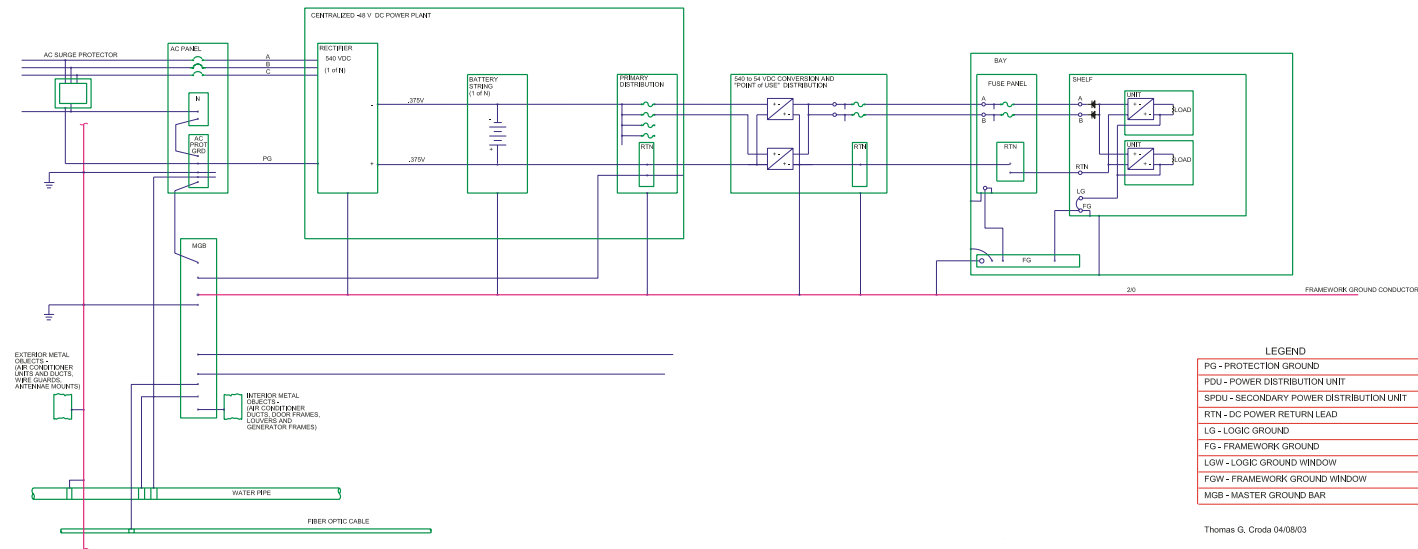
⁶ See also Appendix Q: “Powering the Internet, Datacom Equipment in Telecom Facilities: The Need for a DC Powering Option,” Copyright © 1998 by the Technical Subgroup on Telecommunications Energy Systems of the Power Electronics Society of the Institute of Electrical and Electronics Engineers, Inc.

Figure 5.1.2: 48-VDC system



Thomas G. Croda 04/08/03

Figure 5.1.3: 540-VDC system



Thomas G. Croda 04/08/03

A Preference for DC Power Supply

by Tom Croda

Why are telecommunications switching facilities DC and Internet facilities AC?

The reason telecommunications switching facilities are DC is that they always have been.

The reason that the Internet is AC is that it always was.

There is much inertia resisting change.

In my opinion, DC is a more reliable architecture because it's simple. Multiple paths for redundancy are provided via diodes, not complex switching systems. DC could also be more efficient if all players worked together at a systems level. If functions were served in a redundant manner with a power converter at the point of use, each converter (power supply) could be optimized for the load. Large telecom systems are developed and operated this way.

A large data center could be served by a two-level DC distribution system. Level One would run at 540V DC, the working voltage of 240-cell lead-acid batteries. The primary conversion would be expandable without interruption as the load grew. Energy storage would be done centrally at this level. The primary AC to Level One DC conversion—assuming an AC power source—could be done at 94% efficiency with full transformer isolation. In addition, DC power sources such as fuel cells, certain microturbines, photovoltaics, or wind turbines could input power

directly to the 540-VDC bus, increasing efficiency and helping reduce peak load. Common bus architecture at Level One would operate at 60–80%, eliminating the need for primary converters to operate at 30–40% maximum load.

Level Two would run at the 54-VDC level, the working voltage of most telecom equipment. Conversion to Level Two would occur very close to the load equipment to reduce wire size, and could be 85–90% efficient.

With DC electricity, power can be provided to individual power supplies using diodes from multiple distribution paths. Distributing redundant DC power (at 54 VDC) to small “point of use” DC-to-DC converters on each circuit pack would improve overall system efficiency. In this way, each card has a converter operating at maximum efficiency. When a protective device fails, the total power drain stays the same; the distribution path simply changes.

The primary converters operate within a high efficiency range regardless of the distribution path. The final “point of use” power supplies always operate at the same high-efficiency point.

Much, if not all, of the mechanical system could be supported by VSDs, which are DC devices with the rectifiers removed. Where needed, AC power could be provided via “point-of-use” high-efficiency, modular inverters distributed across the equipment floor near the load equipment.

This architecture could increase efficiency and seriously improve reliability because of its simplicity. It also provides several other advantages in reduced electrical noise and improved transient protection. While such systems are new and will challenge business-as-usual at many conventional equipment providers, most UPS companies already provide chargers and inverters that operate at 540 VDC. The changes described here would only require repackaging.

In terms of efficiency, scalability seems to be the key. In the site efficiency measurements I have taken, that is the most important factor and can be applied now.

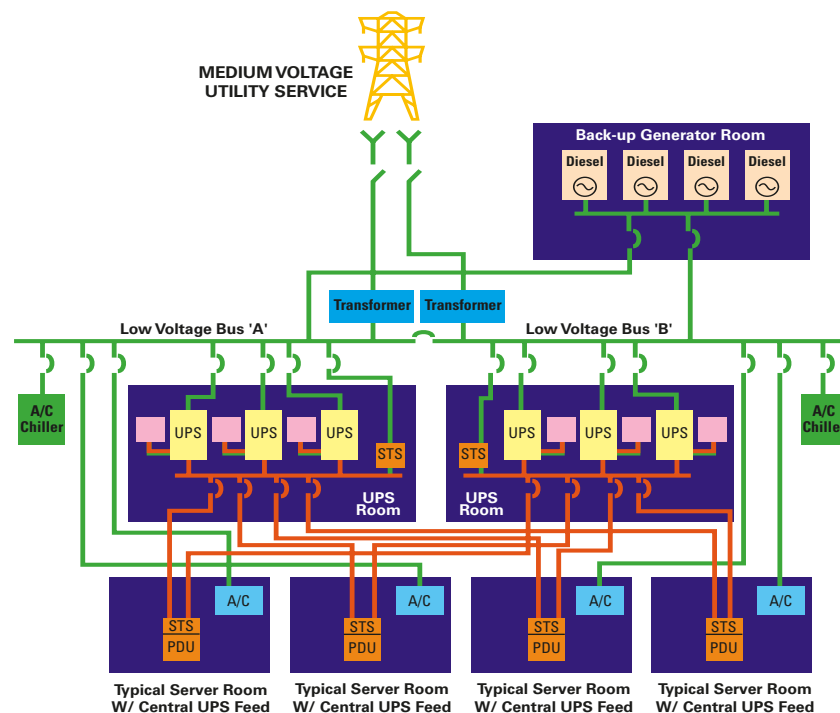
Figure 5.1.2 (on p. 68) is a simple diagram of a 48-VDC conventional system. Figure 5.1.3 (on p. 68) is a quickly adapted version showing a 540-VDC front end. The 540-VDC diagram needs much refinement but it conveys the idea.

5.2 On-site power generation

The primary power supply should be an on-site generation system with, at a minimum, double redundancy, with the grid as backup. Today, in a typical dual-pad, double-redundant system, there are two completely independent buses, each of which have two redundant generators. Because each half of the load runs on two separate redundant systems, however, each generator runs at ~25 percent load. This system is highly inefficient and causes excess heat generation at the generator location, and it is an inefficient use of capital. Figure 5.2.1 (right) shows a typical data center power system used today.

A greatly improved dual-bus system is shown in the circuit diagram in Figure 5.2.2 (on p. 71). This recommended design is scalable, modular, diverse, reliable, and able to load-match. More efficient than the grid, this system uses its waste heat to power a thermal-based cooling system, and reduces overall electrical demand on the system. Rather than using an electrically-driven vapor-compression cycle to drive the heat removal process, a thermal-based cooling system uses the generator's waste heat to drive an absorption, desiccant, or other thermal cooling-cycle technology. This reduces the electricity load for cooling to only the auxiliary loads required to move the necessary air and liquid streams. Moreover, the needs for electrical backup systems and the corresponding losses are also reduced.

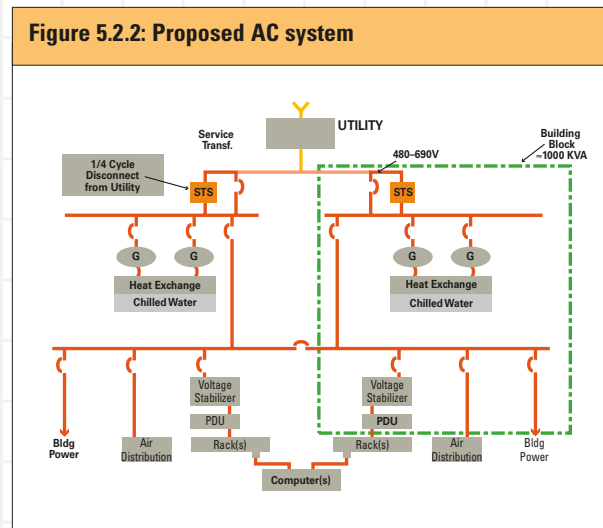
Figure 5.2.1: Typical AC system



This synergy between the data center's requirement for reliable onsite power and its tremendous—but readily reduced—cooling requirement offers a key strategy for reducing overall power consumption.

To simplify the path of power to the computers, the dual-redundant, online UPS systems shown in Figure 5.2.1 are replaced with simpler, self-stabi-

lizing buffer/transient technology systems (e.g., flywheels, new high-power batteries, or ultracapacitors), powered by a clean, reliable onsite power source (e.g., turbines, fuel cells, etc.). Each generator is designed to run at optimal capacity, not the ~25 percent load at which most UPS systems now operate. Ideally, excess power and ancillary services can be sold back to the utility.



Due to the simpler path, the number and capacity of batteries are significantly reduced, as is the need for breakers and switches. The design eliminates 50 percent of the losses that occur in today's systems.

The system is completely scalable and modular. The basic building block is two double-redundant modules, *e.g.*, designed to supply a 100-W/sq.ft. load in a 10,000-sq.-ft. facility with 1 MW of power, while saving at least 500 kW over today's conventional designs. To add capacity as the size of the data center increases (modularly), single modules can be added as necessary. It would be possible to make the individual modules smaller (*e.g.*, 4 × 250 kW). The capacity of the module then determines the minimum increment by which capacity can be increased.

5.3 Interconnect with utility

For the present,⁷ the recommended system should be connected to the grid to ensure reliability and to improve payback. The key to successful connection with the utility is two very fast circuit breakers or static switches on each generating bus⁸ to disconnect the onsite generator quickly and prevent any possible damage associated with reverse fault flows during times of grid failure, when the onsite generator must operate in an "island" mode.

Ideally, an important benefit of interconnection with the utility would be that unused capacity from the redundant generation system's total capacity could be sold back onto the grid. This would allow the flexibility to keep generators running at full load, thus making them optimally efficient (globally more fossil-fuel-efficient). The export of spare power could be an additional revenue source, shortening the payback period of the total investment. Unfortunately, the combination of power export and high-reliability operation is problematic, as discussed in the following recommendation. It may be possible to do this today, however, in an area with an innovative and cooperative distribution utility.

⁷ If in the future there will be very cheap, clean, distributed generation technologies (*e.g.*, cheap fuel cells), then the grid connection becomes unnecessary.

⁸ Several companies make low-voltage "static switches" that operate at power levels from 480 volts to 4000 amps. Others have attempted to do it at medium voltage (5–25kV). S&C Electric claims to be the only one with a true production product. For more information, contact Brad Roberts (see participant directory, pp. 90–99).

5.4 Address barriers to self-generation

Barriers to self-generation include:

- Complex utility contracts that take much time and effort to complete;
- State Public Utility Commission (PUC) requirements;
- Lack of financial incentives in most areas;
- Lack of standard rules for interconnection;
- Cost of maintenance increases by an estimated 1 cent per kWh;
- Utility resistance, requiring case-by-case negotiation;
- Independent cogenerators' expectation of a "take-or-pay" contract (assuming the data center owner/operator is not interested in owning and/or operating these assets itself).

An optimally cost-effective system requires *both* the reliability benefits of standby operation and the energy savings of parallel operation. A critical issue for DG sources is the possibility of "islanding," when a fault in the grid separates a generating source from the rest of the system, creating an electrical "island."⁹ Islanding is essential for


(Continued on next page.)

Recommendation 5.4: Address barriers to self-generation

providing premium customer reliability during grid outages, although current utility practice discourages it.¹⁰ To maintain local power supply during a grid outage, the control system must detect the outage, disconnect from the grid, drop sufficient non-critical load to meet local generation capacity, operate during the outage, and resynchronize with the grid when service returns. Although technically possible, it is difficult under present conditions to design for *both* power export to the grid and for premium reliability by island-mode operation during grid outages. Most distribution utilities will discourage such a configuration. Thus, it is more practical today to design for premium reliability by island-mode operation during grid outages, and for parallel operation under normal conditions without the capacity to export to the grid.¹¹

To help reduce connection and protection costs by making the requirements more predictable, the Institute of Electrical and Electronic Engineers (IEEE) is working to develop a national standard for interconnecting distributed resources with electric power systems, which is expected to be published in 2003.¹² Capturing the potential reliability benefits of onsite generation, without sacrificing the benefits of parallel operation, requires further development of standard practices, in cooperation with distribution utilities. This goal is achievable with existing technology, and this has been demonstrated in practice, including a small number of systems that export power to the grid. However, the capability to *both* export and island increases system complexity and cost, due to the need to avoid system instability in case of a grid outage, and this type of design is generally discouraged by most distribution utilities.¹³

The emergence of net-metering laws or policies in at least 38 states should help to reduce this resistance and to educate utilities about the valuable system benefits of distributed generation.



⁹ An island is "any part of the distribution system, consisting of both generation and load, that operates without interconnection with the bulk power system." Dugan, R. and G. Ball, 1995. *Engineering Handbook for Dispersed Energy Systems on Utility Distribution Systems*. Final Report, Electric Power Research Institute. EPRI TR-105589.

¹⁰ For further discussion of islanding, see *Small Is Profitable* by Amory Lovins *et al.* (Rocky Mountain Institute, 2002), p. 249; www.smallisprofitable.org.

¹¹ National Renewable Energy Laboratory (NREL), 2000. *Making Connections: Case Studies of Interconnection Barriers and their Impact on Distributed Power Projects*, NREL/SR-200-28053.

¹² This standard, IEEE SCC 21 P1547, will include requirements for the performance, operation, testing, safety, and maintenance of DG interconnections.

¹³ A compromise solution might involve separate generation sources, one sized and designed not to export power but to island during a grid outage, and the other designed to export normally and to trip in case of an outage.


Part 6: Operations

There are as many opportunities to improve performance of data centers by correcting the perverse systems governing space, power, and cost as there are by improving power supplies, CPUs, and cooling systems. To improve data center design and operations, incentives must be powerful and relevant, education must be a part of all data center considerations, and disconnected sectors need to work in unison.

Agents all along the value chain need to measure and to pay for the costs of the resources that they demand. The current system of charging users only on the basis of square feet encourages higher density of use and hence higher energy consumption, well beyond the optimum. Current real estate models (design + construction relationships, lease + incentives) generate perverse signals because they do not reflect the true cost of computing power. More importantly, the real estate models don't reflect the true cost of the capital and operating expenses necessary to deliver electricity of the requisite reliability to the server. Thus electricity and capital are used inefficiently. The cost of electric service is the key issue here.

The current real estate model is inappropriate for a data center. A major misconception in space-to-power density ratios is that cost per unit of computation necessarily decreases as power density increases. This is untrue except in certain real estate-oriented pricing structures. As a result of this misconception, many people advocate “compaction,” or cramming more power into smaller and smaller spaces. If properly measured, however, the cost of supplying energy—which includes a huge amount of infrastructure cost and inefficiencies—can be more than \$8,000 per kilowatt. This number represents power *density*, however (see panel discussion in Appendix B), so pricing per square foot and per watt can help more optimally reflect costs and spread power density.




Figure 6a: Educating office owners to charge by the kW as well as by the sq. ft.
How tenant load affects occupancy cost: Class A NYC office in the late 1990s



Ken Brill, Executive Director of The Uptime Institute comments: “What is the ‘true’ cost of providing power, cooling, and space for a cabinet of computer equipment? There are two different ways of charging for occupancy and construction. The current method considers only investment cost and operating cost. In order to calculate the real cost, however, the equation needs to deal with the cost of space separately from the cost of power and cooling. This method points out that minimizing space is stupid; the major cost element is supplying the cooling and power, which leads to radically different conclusions about the economics of further technology compaction.”

There are myriad disconnects between the narrow foci and missions of the individual sector specialists—real estate, facilities, finance, vendors, IT, and end users—and the best interests of the data center as a whole. As Mr. Brill notes, “Real estate professionals get paid to keep costs low; IT guys get paid to make computing space available.” As a result, while minimizing cost, real estate people may unintentionally create unavailability by incentivizing an inherently less reliable infrastructure for which they are not penalized. They are focused on first cost, not operating cost or reliability. Because availability falters, however, the IT professionals are penalized.

The industry will be much better served if it can heal these disconnects and create tight feedbacks between costs and demand. Real estate and IT people need to be better connected to achieve a mature real estate model that serves both. Figure 6a.1 shows one example of a simple tool used for analogous education of landlords and tenants in Manhattan a decade ago, when many leasing brokers and tenants’ agents insisted on wiring and cooling capacity an order of magnitude greater than were actually needed to serve lighting and plug loads—not realizing that extra watts imposed extra costs and thus raised rents. The disconnect is worse with data centers because the added costs are far higher and the price distortion drives high densities that severely compromise the most mission-critical parameter—computer uptime.



Reliability is the most critical element in data center facilities and is the easiest to sell. Therefore, efficiency can not compromise reliability, and success will be facilitated if efficiency is shown to increase reliability. To promote efficiency successfully, it is important to understand its impacts on availability and apply this test to all measures.

The way to capture these opportunities is to make true performance and costs transparent, and get the incentives right. Specific techniques include minimizing idle resources, ensuring there are no excesses and no insufficiencies and “just-in-time computing.”

Recommendation

6.1 Intelligent resource allocation

Work from demand side to supply side to maximize cumulative upstream benefits and to right-size equipment. This is a guiding tenet: figure out what you really want to do, then find the cheapest, most direct way of doing that. Usually, it’s the safest and most reliable way as well.

Recommendation

6.2 Improve information available

- Develop full and disaggregated cost assessments and give them to agents/users/customers all along the supply chain for equipment and electricity. Private and public entities can then make optimized decisions about computing, electricity, and other resources. If people don’t know what something costs and do not have to pay for it, they cannot be expected to optimize its design or use it efficiently.
- Develop methods to calculate lifecycle cost/total cost of ownership. Lifecycle cost should include building and operating costs, renovations, management overhead, maintenance contracts, property taxes, insurance, equipment, energy, software licenses, and anything else that is an expense, including energy (see sidebar: “Lifecycle Costing”). It must also properly reflect uptime by including the cost of downtime; otherwise the optimization will fail.
- Develop success-story case studies of charging on a per-watt basis or other appropriate metric.

Lifecycle Costing

by Greg Kats

Lifecycle costing is never done correctly because those doing it never have complete data. Still, it’s the right tool for many things and it’s advocated here. Some of the standard language is described on the LCA websites listed at right. The list of issues considered does not need to be exhaustive. The elements to focus on are those likely to have significant costs/benefits within the defined timeframe (*e.g.*, 10 years). A lifecycle assessment (LCA) approach to evaluating and integrating benefits and costs associated with sustainable buildings generally involves accounting for all upstream and downstream costs of a particular activity, and integrating them through a consistent application of financial discounting. The result—if data are available—is a current “cradle-to-grave” inventory, impact assessment, and interpretation (*e.g.*, a net present value estimate).

For an extensive international listing of green building evaluation and lifecycle-related tools and programs with related URLs, visit <http://buildlca.rmit.edu.au/links.html>.

For a good overview of international lifecycle development, see “Evolution and Development of the Conceptual Framework and Methodology of Life-Cycle Assessment,” SETAC Press, January 1998. Available as an addendum to a “Life-Cycle Impact Assessment: The State-of-the-Art.” See www.setac.org.

6.3 Align incentives with desired performance

- Eliminate today's perverse incentive structures by aligning market incentives with *desired performance*. Instead of charging on a per-square-foot basis, select from a diverse menu of interrelated incentives: per watt, per unit of power density, per teraflop, etc.—whatever metrics are practical, efficient, and indicative of key costs and benefits.
- Use performance-based fees to provide incentives to encourage design teams to create buildings and equipment that are optimally efficient. Performance-based fees reward the team from the savings generated by better design.
- To increase reliability, the pain of failures must be shared. One sector should not be penalized so that other sectors might be rewarded; all should share in successes and failures in terms of energy consumption. One way to achieve this is to create service agreements to link IT and real estate functions, and deal with:
 - Crises;
 - Price and economic motivation;
 - Risk management concerns; and
 - Service level agreements.

6.4 Benchmarking

Gathering and benchmarking operating data about computing facilities and data centers is essential, and is a key recommendation of the charrette. Feedback on costs is essential both for operations (short run) and planning (long run) of data flow and processing capacity. The data must be globally available, transparent, and translatable across boundaries. The collection and distribution of these data may well be most-appropriately web-based.

To facilitate understanding of data center power and costs, comprehensible and useful metrics must be developed and benchmarked. Some of these metrics are not yet in use; others are easily calculated from existing data.¹ Recommended metrics include:

- Metric of computational output²—kW per unit of computational output;
- kW per rack equivalent—allows tracking of “packing factor”;
- UPS efficiency or losses—ratio of total kW in to UPS power output, kW of HVAC/kW of UPS;
- Plug-process load W—W/ft² nameplate energy labeling for peak, end use, idle, power supply efficiency;
- Total kW demand per kW provided to the servers (a measure of parasitic power demand) or to all IT equipment, or the ratio of electrical computer equipment load to the total building or data center load (this would be a measure of the infrastructural energy efficiency);³
- Cooling—kW/ton, ft²/ton, unit of cooling per unit of data-processing;
- Air recirculation—cfm/ft², W/cfm, air changes per hour in computer room;
- Power transformer efficiency—percent efficient;
- Lighting—W/ft² as used (net of any control savings); and
- Effective air infiltration or leakage area effect.

¹ Possible data sources include: Uptime Institute, 7X24 Exchange, PG&E, LBNL (see <http://datacenters.lbl.gov>), utility audits, Energy Star Data Center Rating (in process), and measurements from commissioning reports that capture 100% load data.

² Review comment from Eng Lock Lee: “Dividing the energy by MIPS yields joules/instruction, *i.e.*, (joules/sec) / (i/sec). This is one measure of the energy efficiency of the CPU; one could also use I/O or flops or some other metric. I wonder if someone has done the exercise and worked out the net efficiency of the CPU and supporting devices on typical problem types, *e.g.*, to simulate collision of protons it takes 100,000 Btu of electrical energy to the semi-conductor devices, and also 50,000 Btu of HVAC support.”

³ Review comment from Bernard Aebischer: “In the nineties, energy efficiency of the central infrastructure of a group of computer centers in Switzerland was benchmarked using the ratio of energy used by the computers divided by total electricity used in the computer center. The Canton of Geneva is proposing to use this coefficient ‘C1’ in the construction authorization process for new data centers and for defining targets in a voluntary commitment process applicable for new and existing data centers.”

6.5 Write more efficient code

Eliminate “bloatware” and make code that allows chips to scale up and down. Bloatware refers to the drive to make continually larger, fancier applications⁴ that results in users’ continually upgrading equipment to run those applications.

Recommendation

6.6 Submetering

Submetering end-uses allows real time feedback and adjustments to reflect real costs. Its use is important for Energy Star building level and LEED ratings. Effective submetering requires accurate, distributed sensors. Today, very little submetering is done. More conducive utility rules, and utility rebates based on submetering, would encourage this practice.

Recommendation

6.7 Measurement and verification (M&V)

M&V capabilities continue to improve rapidly while costs decline, allowing more cost-effective real-time monitoring and management of energy and buildings systems to increase systems performance (including energy savings) improve system reliability, and reduce mean time to failure.

⁴ Example: how much better does the typical word processing software run now—for the small percentage of its functionality you actually use—than it did 10 years ago?

Measurement and Verification

“You cannot manage what you do not measure.”

—Jack Welch, CEO of General Electric

Working with industry to overcome existing barriers to efficiency, the U.S. Department of Energy developed a consensus approach to measuring and verifying efficiency investments. The International Performance Measurement and Verification Protocol (IPMVP) was first published in 1996. North America’s energy service companies have adopted the IPMVP as the industry standard approach to measurement and verification (M&V).

The International Performance Measurement and Verification Protocol (MVP) provides an overview of current best practice techniques available for verifying results of energy efficiency, water efficiency, and renewable energy projects. It may also be used by facility operators to assess and improve facility performance. Energy conservation measures covered include: fuel saving measures, water efficiency measures, load shifting and energy reductions through installation or retrofit of equipment, and/or modification of operating procedures.

Simply put, the purpose of the IPMVP is to increase investment in energy efficiency and renewable energy. When firms invest in energy efficiency, their exec-

utives naturally want to know how much they have saved and how long their savings will last. The determination of energy savings requires both accurate measurement and replicable methodology, known as a measurement and verification protocol.

The key to unlocking the enormous potential for energy efficiency worldwide is securing financing. This requires confidence that energy efficiency investments will result in a savings stream sufficient to make debt payments. Measurement and verification practices allow project performance risks to be understood, managed, and allocated among the parties.

The Protocol:

- Provides a common set of terms and establishes methods which can be used in energy performance contracts.
- Defines broad techniques for determining savings.
- Applies to a variety of facilities.
- Provides internationally accepted, impartial, and reliable outline procedures.
- Provides a comprehensive approach to building indoor environmental quality issues.
- Creates a living document.

For information on M&V see www.ipmvp.org.

6.8 Continuous commissioning

Implement and maintain a comprehensive “best practices” and continuous maintenance system. Continuous commissioning can:⁵

- optimize operation of existing systems;
- improve building comfort within the capabilities of the installed system;
- reduce building energy cost;
- reduce operational and maintenance costs;
- help to ensure continuous optimal operation for years to come;
- improve technical knowledge of operating personnel; and
- usually pay back in less than two years.

The continuous commissioning process:

- investigates and documents the condition of the mechanical systems;
- solves existing problems in the building within the capabilities of the installed system;
- optimizes building energy systems and formalizes operational procedures;
- measures and documents the energy savings and comfort improvements; and
- provides ongoing monitoring of system operation.

6.9 Create self-diagnosing/healing systems

- **Computing:** Systems sense faults, make corrections, and self-allocate hardware resources to meet demand. Smart systems would be able to draw from all data center resources for free RAM, available CPU capacity, storage, etc.
- **HVAC:** Use RAIS (Redundant Array of Inexpensive Sensors) to measure thermal conditions throughout the data center and trigger appropriate dynamic responses to changing conditions.

6.10 Virtual servers

A large mainframe that hosts many virtual servers appears to the outside world to be many different servers. To the owner it’s one big machine. There are many advantages to this arrangement, and it solves many of today’s issues. Operators are likely to move more towards this approach in the future.

6.11 Optimization tools

Create distributed planning, forecasting, and design tools for data center end users and designers to provide price signals that reflect true costs, as well as dynamic tools that simplify design and construction of efficient devices.

6.12 Miscellaneous

- Apply the experience and knowledge gained from energy demand side management (DSM) programs for the cost-effective management of computer resources.
- Increase modularity of all components, especially large equipment in data centers.
- Minimize administrative burdens and transaction costs.

⁵ Source: Energy Systems Laboratory, Texas A&M University Energy Systems Opportunity Assessment. See <http://energy.opp.psu.edu/engy/CCommiss/CComHome.htm>.

6.13 Education, outreach and training

- To educate properly about energy efficiency, data center users, developers, and owners need to define the end users who need education, and define their motivations (IT, facilities management, consulting engineers, enterprise level CIO, CTO, IT consultant, etc.).
- Knowing this, create customized curriculum and delivery mechanisms for each group.
- Create “best practices” manual based on existing technologies.
- Prepare and disseminate case studies.
- Define commissioning program.
- Best practice center/testing lab/education center—a location where various technologies can be tested, tinkered with, and showcased; similar to the existing Hewlett-Packard data center test lab.
- Find collaborative funding for and operation of educational programs (industry sponsors, DOE, OIT, EPA, utilities, CEC, DOD, etc.).
- Create a market for reliability through energy efficiency.
- Work with consultants to publish/recommend.

Mechanisms for the delivery of education include:

- Utility DSM programs such as audits and incentives funded by a public goods charge. These might be offered free or at a negotiated cost based on energy savings.
- Create a “Data Center for Excellence” program, possibly in affiliation with groups such as the USGBC, EPRI, or the Silicon Valley Manufacturing Group.
- Organizations with related interests (*e.g.*, ASHRAE, 7X24, Uptime Institute, AFCOM, SHARE, BOMA, IFMA, etc.)
- Create a “Data Centers for the 21st Century” similar to the existing “Labs for the 21st Century.” This could include a LEED-type rating system for data centers.
- The Collaborative for High-Performance Schools was suggested as a model for data centers. It includes nonprofits, governments, and the private sector.

6.14 Demonstrations

IT people are risk-averse; they need to be shown how well low-power data centers and their components can perform, how secure they can be, and that risks are likely to decrease with these types of data centers. These recommendations may be best proven and demonstrated via a pilot project data center or some other type of showcase project—which should be in a green building.



6.15 Energy Star and LEED ratings

Creating standards to measure efficiency provides incentives to improve efficiency. Standards for efficiency and for comparing data center energy use prompt companies to design data centers to achieve high performance. Efficiency requirements encourage manufacturers to modify their systems to live up to at least the minimum requirements. These are not hypothetical statements. The Energy Star program and LEED system for “green” buildings (see sidebars) are demonstrating that these things actually do happen.

Establishing Energy Star ratings for servers and data center cooling systems will build consensus on both what the real heat loads and efficiencies of components are and which system architectures offer optimum energy efficiency. With industry buy-in, Energy Star ratings can be established relatively quickly and implemented voluntarily.

Implementation

Implementation: Several things are needed before Energy Star ratings can be applied to servers and data centers. These include:

- numerical, quantifiable statements about energy usage;
- good metrics (flops/W, calculations/cycle, etc.); and
- good baselines (find a good model somewhere).

As a first step, consider creating an Energy Star rating that focuses only on power supply and fan efficiency. This avoids the problem of defining performance metrics because these measures are independent of the processor that is used in the server.

LEED



The U.S. Green Building Council (USGBC) is a leading national organization that promotes the construction of and creates standards for energy- and resource-efficient buildings. LEED™ (Leadership in Energy and Environmental Design) is a voluntary, consensus-based national standard developed by the USGBC for developing high-performance, sustainable buildings. It provides a comprehensive framework for assessing building performance and meeting sustainability goals. Based on well-founded scientific standards, LEED emphasizes state-of-the-art strategies for sustainable site development, water efficiency, energy efficiency, materials selection, and indoor environmental quality. LEED recognizes achievements and promotes expertise in green building and offers project certification, professional accreditation, training, and practical resources. Council members work together to develop LEED “products” and resources, policy guidance, and educational and marketing tools that support the adoption of sustainable building. About 8% of all new commercial buildings in the United States in 2003 are being LEED-certified, and market demand for such certification is strongly influencing designers’ skill sets and marketing. For more information and to download the free guideline matrix, see www.usgbc.org.

6.16 Create an independent organization to provide testing, experimentation, education, and demonstrations

This could be a nonprofit organization run by a consortium of utilities, manufacturers, data center operators, state energy research agencies, and other interested parties. It could take the shape of a subscription service provider. It could grow out of existing trade organizations that currently have a narrower focus. If necessary, it should be created and jump-started by state energy efficiency agencies that manage public goods fees.

Many functions that such an organization could provide are discussed in this report. These include:

- Performing essential benchmarking upon which to base all other actions. Gathering and benchmarking information about computing facilities and data centers is essential, and is a key recommendation of this charrette;

- Developing broad-based requirements and standards for industry declaration of power supply performance statistics, including efficiency versus load curves, to show part-load performance;
- Organizing collaborative funding for and operation of educational programs;

(Continued on next page.)

ENERGY STAR



ENERGY STAR is a government-supported program that promotes energy efficiency.

In 1992, the U.S. Environmental Protection Agency (EPA) introduced ENERGY STAR as a voluntary labeling program to identify and promote energy-efficient products with low greenhouse gas emissions. Computers and monitors were the first labeled products. Between 1992 and 1995, EPA expanded the label to additional office equipment and residential heating and cooling equipment. In 1996, the EPA partnered with the U.S. Department of Energy (DOE) for particular product categories. The ENERGY STAR label is now on major appliances, office equipment, lighting, home electronic devices, and other items. The EPA has also extend-

ed the label to cover new homes and commercial and industrial buildings.

Through its partnerships with more than 7,000 private and public sector organizations, the ENERGY STAR program delivers the technical information and tools that organizations and consumers need to choose energy-efficient products and best management practices. Over the past decade, ENERGY STAR has been a driving force behind the widespread use of such technological innovations as LED traffic lights, efficient fluorescent lighting, power management systems for office equipment, and low standby energy use.

EPA provides an innovative energy performance rating system that businesses have used for more than

10,000 buildings across the country. EPA recognizes top performing buildings with the ENERGY STAR. Because a strategic approach to energy management can produce twice the savings—for the bottom line and the environment—as typical approaches, the EPA's ENERGY STAR partnership offers a proven energy management strategy that helps in measuring current energy performance, setting goals, tracking savings, and rewarding improvements.

Results are adding up. Last year alone, Americans, with the help of ENERGY STAR, saved enough energy to power 10 million homes and avoid greenhouse gas emissions from 12 million cars—all while saving \$6 billion.

For more information see www.energystar.gov.

Recommendation 6.16:

Create an independent organization to provide testing, experimentation, education, and demonstrations

- Showcasing new technologies to IT professionals and consumers to prove that such things as off-server disks are fast and reliable, and that system security issues can be addressed;
- Making it clear to power supply, processor, HVAC equipment, and other manufacturers and suppliers that power and efficiency are major concerns to data center owners/operators;
- Helping create LEED standards for data centers and Energy Star ratings for servers and data centers;
- Defining operating envelopes and establishing environmental standards for mechanical and electrical systems;
- Developing full and disaggregated cost assessments for equipment and electricity. Private and public entities can then make optimized decisions about computing, electricity, and other resources;
- Overseeing the development of standards for connections to electrical and liquid cooling systems. Lack of standardization is the principal barrier to the widespread adoption of both blade servers and liquid cooling; and
- Addressing barriers to self-generation, interconnection, and power export.

Uptime Institute's Mission-Critical Product Certification (MCPC) Program



The Uptime Institute, Inc. is establishing a product verification and testing program for mission-critical equipment. The program's mission is to accelerate the adoption of new technology and encourage the enhancement of existing technology in mission-critical facility products. The program does this by reducing performance and reliability uncertainties that end users and those in the engineering community face when making purchasing and deployment decisions.

The program will comply with American National Standards Institute (ANSI) guidelines for testing and certifying bodies. The principal tenets of these guidelines are neutrality, independence, openness, and free access to information by all affected parties, including product manufacturers, end users, consulting engineers, and other interested individuals and companies.

Specific Program goals include:

- Developing and maintaining industry accepted standard measurement and evaluation methods applicable to different classes of products used in mission-critical facilities. These criteria will be established utilizing the collective experience and wisdom of recognized industry leaders (end-users, consulting engineers, and subject matter experts).
- Conducting independent, rigorous, and verifiable performance testing of similar products from multiple vendors.
- Publishing test reports, allowing prospective purchasers to make decisions based on standard measurement and evaluation methods.

Conclusion

A handful of broad conclusions emerged from this charrette. They include:

- Low power data centers (with low power servers and CPUs) can be created using current technology, as can all the individual items recommended in this report: less-energy-intensive components (batteries, CPUs, chillers, fans, etc.), conductive cooling (water), natural ventilation, redesigned servers (with fans, OS, etc. off the rack), etc.
- Gathering and benchmarking data about computing facilities and data centers is essential. The data must be globally available, transparent, and translatable across boundaries. The collection and distribution of the data will probably best be web-based.
- All individuals involved in the planning, designing, siting, construction, operation, and maintenance of data centers need to share goals and information and any pain throughout all stages of the process. One sector should not be penalized so that other sectors might be rewarded; all should share in energy efficiency's successes and failures so that all actors can learn quickly and continuously improve industry practice.

- A significant amount of education is required for the creation of more efficient data centers. This needs to start with the simple fact that data centers are necessary for modern life, and that current computing systems, through deficient design, threaten the vital information that they process.
- Efforts at data center redesign need to be realistic, scalable, geographically repeatable, economically sensible, and as transparent as possible.

So what are the next steps in the evolution of data centers? Many of them are outlined in this report. Certainly, the steps taken by Lawrence Berkeley National Laboratory have indicated that there is both interest and room for a great deal of improvement. LBNL is already working with CEC and NYSERDA on various pilot projects, notably a roadmapping project and an air management project, aimed at energy consumption reduction. Additionally, Pacific Gas & Electric plans to share a variant of this report with developers, designers, and architects involved in energy-efficient design via the Energy Design Resources website.¹

¹ www.energydesignresources.com.

One goal of this report is to stimulate further examination of the various components of data centers and the energy they consume. At the same time, the report points out why these components must be designed and combined in an integrated—whole-systems—fashion.

The charrette results clearly pointed out how quickly the value of saving one watt compounds throughout the total data center system. We detailed a reduction of 83.5 percent in the computing equipment itself. This translated into a 94 percent reduction in all the other building system loads that support the equipment loads. This illustrates how the savings in one system cascades into numerous related systems. Additionally, looking only at energy consumption does not reveal other operational costs, such as human costs and the lost revenue from downtime and unreliable performance and the simple costs of maintaining the systems. Finally, in the case of data centers, efficient design massively reduces the quantity of material resources needed to provide computing services.

The Internet has become an increasingly important factor in the national and global economy. At this charrette we were able to take advantage of the current business slowdown to step back and critically examine current practices. We can expect aggressive growth of Internet-related facilities to resume. When that happens, we hope that the ideas developed at this charrette and presented by our report will help to ensure orderly, profitable, and environmentally responsible growth.

How quickly will the data center of the future be realized? We don't know, but the early-21st-century lull in the economy and the bursting of the late-1990s technology bubble have provided all who work with data centers, computers, and high-tech real estate an important chance to do data centers right the second time.

We hope that readers will use this report as inspiration to challenge conventional designs of buildings, servers, CPUs, and support systems. But most importantly, we hope you will use it to challenge conventional thinking about energy consumption, and how we design and build systems around bits and bytes.

N1's computing-on-demand to drive network services

by Greg Papadopoulos

The next big thing for Sun, and the industry, is the N1 architecture, which will simply treat the network as a computer.

Since we entered the business 20 years ago, the definition of a system has remained constant amid rapid refinements. The components have always included processors, disks, memory, and network I/O.

For the next 20 years, things will look much different. In the N1 architecture, the components will include computers, storage systems, and IP networks.

In the current view of computer systems, the units of work are known as processes. Moving forward, the units of work will be web services. N1 is designed to create a single pool of resources that can be dynamically provisioned to meet the needs of a whole list of services. Yousef Khalidi, a Sun Microsystems engineer and principle architect behind N1, said that the idea behind N1 is to match resources to services on the fly. "Whenever demand for a service goes up or down, the N1 architecture adjusts to it automatically," he said.

One key result is that change-management will be automated, complexity reduced, resources better utilized, and total cost of ownership lowered.

N1 is an open architecture that will provide a means to virtualize the elements of the network—"the servers, the storage, even the cabling"—so that they can be easily managed. Further, N1's dynamic resource allocation means redundancy and high availability are already built in and need not be added as an afterthought.

Systems based on N1 will be designed according to its guidelines that redefine how systems resources "processing, persistence, communications" are used and organized. With N1, computers don't just attach to networks, they are built from networks. This shift enables radically higher-scale 10,000-plus processors, exabytes of storage, terabits of bandwidth, and millions of IP connections, all of which will be imperative as we move forward.

N1 represents an extension of grid computing and the whole utility model for delivering services on demand.



Rocky Mountain Institute (www.rmi.org), founded by Amory and Hunter Lovins in 1982, is an independent, entrepreneurial, nonpartisan, nonprofit applied research center.

Its ~50 staff foster the efficient and restorative use of natural, human, and other capital to help make the world secure, just, prosperous, and life sustaining. The Institute's ~\$6-million annual budget comes roughly half from programmatic enterprise earnings, chiefly private-sector consultancy, and half from grants and donations.

RMI is known worldwide for its work in advanced resource productivity, business innovations related to natural capitalism, and highly original transdisciplinary syntheses at the nexus of energy, resources, environment, development, and security.

RMI Staff and Consultant team

William Browning,

principal and founder of RMI's Green Development Services
970.927.7304 / browning@rmi.org

Cameron Burns, RMI writer and editor

970.927.7338 / cameron@rmi.org

Odd-Even Bustnes, special aide to Amory Lovins

970.927.7343 / oebustnes@rmi.org

Craig Collins, RMI Energy & Resources Services

craigcollins@rmi.org

Huston Eubank, AIA, principal and charrette co-organizer,

a member of RMI's Green Development Services
970.927.7317 / huston@rmi.org

Corey Griffin, Former RMI/GDS Konheim Intern

cgriffin@uclink.berkeley.edu

Marty Hagen, RMI information systems manager

970.927.7314 / mhagen@rmi.org

Onno Koelman,

RMI's Mineral Acquisitions Partnership Fellowship intern
970.927.7301 / okoelman@rmi.org

Amory Lovins, Convener, chief executive officer of RMI

970.927.3851 / ablovins@rmi.org

Jenifer Seal,

principal, a member of RMI's Green Development Services team
970.927.7322 / jenseal@rmi.org

Joel Swisher, PhD, PE, principal and team leader of RMI's Energy &

Resources Services 303.245.1003 / jswisher@rmi.org

Tom Watanabe,

charrette co-organizer, RMI business development director
510.526.5644 / tomwatanabe@rmi.org

Devra Bachrach

Natural Resources Defense Council (NRDC)
71 Stevenson St., #1825
San Francisco, CA 94105
415-777-0220
dbachrach@nrdc.org

Dick Bourne

Davis Energy Group, Inc.
123 C St.
Davis, CA 95616
530.753.1100
dbourne@davisenergy.com

Malcolm Lewis

Constructive Technologies Group (CTG)
16 Technology Dr., Suite 109
Irvine, CA 92618
949-790-0010
mlewis@ctg-net.com

Piers Heath

Battle McCarthy
Dog and Duck Yard
Princeton St.
London WC1R 4BH
+44 (0)20 7 440 8261
pheath@battlemccarthy.com

Gregory Kats

Capital E
1828 L. St. NW
Suite 1046
Washington, DC 20003
202-463-8469
www.cap-e.com
gkats@cap-e.com

Ron Perkins

Supersymmetry USA, Inc.
4509 Evergreen Forest Lane
Navasota, TX 77868
936-894-2819
ronaldperkins@msn.com

Peter Rumsey

Rumsey Engineers, Inc.
99 Linden St.
Oakland, CA 94607-2508
510-663-2070
prumsey@rumseyengineers.com

Dale Sartor

Environmental Energy Technologies Division
Building Technologies Department
Lawrence Berkeley Nat. Lab (LBNL)
MS 90- 3136
Berkeley, CA 94720
510-486-5988
http://eetd.lbl.gov/btp/ateam/ateamweb/tea_m/dale.htm
DASartor@lbl.gov

J.B. Straubel

Volacom
1509 East Garfield Ave.
Glendale, CA 91205
650-996-4355
straubel@stanfordalumni.org

Geoff Wood

Profile Composites, Inc.
1416 Lands End Rd.
Sidney, BC V8L 5K1
250-655-7142
GMWood@aol.com

Volunteers

Gautam Barua

The Natural Step
116 New Montgomery St.
Suite 800
San Francisco, CA 94105
415-531-5317
gbarua@naturalstep.org

Michael Daish

616 Pomeroy Ave.
Santa Clara Ca 95051
408-455-3145
mdaish@email.com

Ann Hushagen

Oregon Office of Energy
625 Marion St. N.E.
Salem, OR 97301
503-373-7804
Ann.Hushagen@state.or.us

Greg McBeth

piguy@Stanford.EDU

Cody Taylor

P.O. Box 13698
Stanford CA 94309
cody1984@stanford.edu

Devra Bachrach is a scientist with NRDC's energy program. Based in the San Francisco office, she works on promoting energy efficiency and renewable energy resources through research and advocacy in local, state, and regional forums. Prior to joining NRDC in 2002, Devra worked on energy policy at the California Energy Commission during the 2001 energy crisis, and conducted research on renewable energy and long-term power contracts at Lawrence Berkeley National Laboratory. She received her master's degree in energy and resources from the University of California at Berkeley, and a bachelor of science degree in engineering from UC Berkeley.

Richard Bourne, PE, with the Davis Energy Group, earned a master's degree in mechanical engineering from Stanford in 1968 and specializes in the design, development, and performance prediction of energy-efficient projects and components for buildings. A former Associate Professor at the University of Nebraska, Mr. Bourne has presented more than 150 special lectures, workshops, and technical papers on energy subjects since 1974. He served as Chairman of ASHRAE's radiant heating/cooling technical committee, and holds 14 U.S. patents. Mr. Bourne is co-founder of Davis Energy Group, served as President from 1981 to 1997, and currently leads their team in development of emerging technologies.

William D. Browning is a Principal with Rocky Mountain Institute's Green Development Services team. Browning has a bachelor's degree in environmental design from the University of Colorado, and a master's degree in science in real estate development from the Massachusetts Institute of Technology. In 1991, Browning founded Rocky Mountain Institute's Green Development Services, a program that researches and provides consulting on environmentally-responsive real estate development. Green Development Services was awarded the 1999 President's Council for Sustainable Development/Renew America Prize. Browning has given numerous lectures and workshops for architectural and environmental organizations throughout the world. He is a coauthor of *Green Development: Integrating Ecology and Real Estate*, a companion CD-ROM, *Green Developments: A Primer on Sustainable Building*, and *Greening the Building and the Bottom Line*. In 2001 he was selected as an honorary member of the American Institute of Architects, and in 2002 was brought onto the national steering group

for the AIA Committee on the Environment. He serves on numerous boards, including the Trust for Public Land and USGBC, and is a principal in the development of Haymount, a new town in Virginia.

Cameron Burns is Staff Editor at Rocky Mountain Institute. A longtime Colorado-based writer, photographer, designer, and communications specialist, Burns received a bachelor's degree in environmental design from the University of Colorado, with an emphasis on architecture and urban design. After graduating, he took up writing, photography and design full-time. As a writer, he has authored or co-authored 17 books, and several thousand magazine, newspaper, journal, and Internet articles. A former environment writer/correspondent for the *Aspen Times* and the *Rocky Mountain News*, much of his writing has focused on energy- and resource-related issues. He has written for many energy magazines, including *Solar Today*, *Nikkei Ecology* (Japan), *Greenbiz.com*, and *EnergyPulse.net*.

Odd-Even Bustnes is Special Aide to Rocky Mountain Institute CEO Amory Lovins and a consultant advising selected RMI clients in the energy and commercial-industrial sectors. Before joining RMI, Mr. Bustnes was an associate with McKinsey & Company, Inc., a Washington, DC consulting firm. He worked there with middle and upper management of major corporations on operational and strategic issues, principally in energy, telecommunications, and private equity. In 1998, Mr. Bustnes analyzed companies in the shipping industry for the Union Bank of Norway Securities, in Oslo. In 1997, as an analyst for UNICEF in Lima and Cuzco, he evaluated water systems in 37 Andean communities. Mr. Bustnes received his MSc in chemical engineering from University College, Oxford University, in 2002. He earned his 1999 MPA in economics from Princeton University's Woodrow Wilson School. He graduated with high honors from Dartmouth College in 1994 with a BA in engineering and government. He speaks fluent Norwegian, English, and Spanish. He achieved the rank of corporal in the Norwegian Army Special Forces and graduated among the top 10 in an elite paratrooper platoon. Mr. Bustnes was a member of the Norwegian Olympic Rowing Team, finishing eighth in the straight fours event in the 1996 Games. He has taught ice and rock climbing and cross-country skiing in Norway. He is a member of the American Council for Renewable Energy and the American Alpine Club.

Huston Eubank is a Principal with Rocky Mountain Institute's Green Development Services team. He is a registered architect with 30 years' diverse experience, both inside and outside the construction industry. His recent consulting projects for GDS have included Lucasfilm's Letterman Digital Center, the Boston Convention Center, Bulmers (the world's largest maker of hard cider), Greening the California State Capitol, and environmental design for schools in Brazil. Prior to joining RMI, he was Director of Building Futures Services at Gottfried Technology, Inc., where he helped develop and implement energy-efficient and environmentally-responsible solutions for new commercial construction. Earlier, as the "Earth Smart Ambassador" for Enron/Portland General Electric in Portland, OR, he managed an innovative energy-efficiency utility program. He was a founding member of the Oregon Natural Step network, liaison to the NW Regional Council of the President's Council on Sustainability, Chair of the Los Angeles CSI Technical and Environmental Committees, Chair of the Oahu Citizen's Advisory Committee on Coastal Zone Management, on the Board of Envirosense, and an advisor for the Architecture+Energy Award program. He is LEED-accredited and the recipient of a bachelor of architecture degree from Cornell University. He has served as a senior project architect at Gensler, a certified construction specifier, a developer, a contractor, a U.S. Navy officer, and the proprietor of his own firms.

Marty Hagen, Team Leader and Information Systems Manager at Rocky Mountain Institute, graduated *magna cum laude* with a BS from California Polytechnic State University. He was a senior field technician at CompuCom, in California, where he provided Macintosh and Windows 95 technical support for Adobe Systems employees. At Kenetech Windpower, he simulated structural loads on wind turbines and developed a computer program to predict wind turbine energy production. With NASA, he conducted aerodynamics, turbulence, and acoustics measurements of helicopter rotors and worked on other aspects of rotorcraft aerodynamics. In the U.S. Army, he repaired anti-tank guided missile systems, worked in nuclear, biological, and chemical defense, installation security, and counter terrorism, and received the Army Commendation Medal. His technical papers have been published in the *Journal of the American Helicopter Society* and in NASA journals.

Piers Heath is Principal Mechanical Project Engineer and leads Battle McCarthy's Environmental Simulation Group. Battle McCarthy Consulting Engineers (a London-based office) leads holistic engineering design teams covering all building and infrastructure design with priority of integrating the design process. Mr. Heath has worked as a building services engineer for over 18 years. He has worked in a number of major and internationally-renowned engineering firms including Ove Arup & Partners and Brown and Root. He specializes in the design of low-energy environmentally-responsive buildings and has extensive hands-on knowledge of very sophisticated computer modeling techniques, including CFD and Dynamic Thermal Modeling. He has previously been involved in a number of major buildings including the refurbishment/change of use for the Pearl Insurance building in Holborn to a five star hotel, which included careful design of systems acceptable to English Heritage using CFD modeling as a proving method. He also worked on a natural ventilation scheme for a 19-story office tower in Kenya, using TAS for modeling. The building was designed without air conditioning and with a minimal dependence on mechanically-controlled components.

Joanie Henderson is currently working in the Commercial and Industrial Services group of Rocky Mountain Institute. Recent projects have been the creation and presentation of Innovation Labs for Royal Dutch Shell. She has worked with Global Partners for Development and World Neighbors. She has also worked with Habitat for Humanity (H4H), educating the members on issues of environmental responsibility, ranging from the compilation of alternative building materials and local availability, to designing for energy efficiency and the principles of passive solar. She has experience teaching and installing small-scale energy sources including biomass, wind, photovoltaic, solar-thermal, and hydro. Her experience also includes working with distributed generation, independent power producers, and co-generation facilities.

Gregory Kats is cofounder and Principal with Capital E (www.cape.com), a national provider of integrated intelligence, strategic consulting, technology assessment, and investment advisory services in the distributed energy industry. Capital E clients include Fortune 100 companies, venture capital firms, clean energy technology start-ups, and

public sector clients. Mr. Kats served as the Director of Financing for the \$1.2 billion dollar Office of Energy Efficiency and Renewable Energy at the U.S. Department of Energy. He initiated and managed the development of large-scale, innovative financing initiatives to support clean energy projects, including a national green power insurance/financing initiative that the *Financial Times* described as "remarkably high leverage." Mr. Kats co-founded and, from 1995 to 2001, served as Chairman of the International Performance Measurement & Verification Protocol (www.ipmvp.org), which involves hundreds of corporations and financial and energy-related institutions. The IPMVP is now the de-facto U.S. standard, has served as a technical basis for over \$3 billion in comprehensive building energy upgrades, and has been translated into 10 languages.

Onno Koelman graduated from Stanford University with a bachelor's degree in mechanical engineering. He has worked as an efficiency expert for a waste collection company, and won a MAP fellowship to do research at Rocky Mountain Institute on the subjects of biomimicry, biophilia, and also an energy resource plan for the city of San Francisco.

Malcolm Lewis, PE, is President and Founder of Constructive Technologies Group. Dr. Lewis is a consulting engineer who specializes in mechanical, electrical, and energy systems for buildings and industrial processes. He has a vast amount of specialized experience in the introduction of innovative building technologies and design processes. These technologies and processes include energy efficiency, sustainable building design, daylighting, thermal energy storage, and cogeneration facilities. Dr. Lewis has over 25 years' of experience in engineering design and the analysis of energy-using systems in buildings. He is the engineer of record for hundreds of new construction and renovation projects for both public- and private-sector facilities. These facilities total over 25 million square feet. Dr. Lewis has been responsible for the design of energy-efficient facilities including central plants with thermal energy storage up to 20,000 ton-hours, cogeneration, and power generation facilities up to 2.5 megawatts, and buildings up to 250,000-square-foot that incorporate daylighting and high-efficiency HVAC and lighting systems. His past project work has included such diverse technologies as fuel cells, active and passive solar heating and cooling, wind power, and photovoltaic power. Dr. Lewis has served as peer reviewer for

numerous energy-conscious design projects throughout the United States and abroad. He has been an energy consultant to the State of California Office of Energy Assessments, The World Bank, Southern California Edison Co., Los Angeles Department of Water and Power, and Southern California Gas Co.

Amory Lovins is cofounder and CEO of Rocky Mountain Institute (www.rmi.org), a 20-year-old, ~50-person, independent, entrepreneurial, nonprofit applied research center in Old Snowmass, Colorado. RMI fosters the efficient and restorative use of natural and human capital to create a secure, prosperous, and life-sustaining world. Mr. Lovins also founded and chairs RMI's fourth for-profit spinoff, Hypercar, Inc. (www.hypercar.com), and cofounded its third, E SOURCE (www.esource.com), which was sold to the *Financial Times* group in 1999. A consultant physicist educated at Harvard and Oxford, he has received an Oxford MA (by virtue of being a don), eight honorary doctorates, a MacArthur Fellowship, the Heinz, Lindbergh, World Technology, and Hero for the Planet Awards, the Happold Medal, and the Nissan, Mitchell, "Alternative Nobel," Shingo, and Onassis Prizes; held visiting academic chairs; briefed 16 heads of state; published 28 books and several hundred papers; and consulted for scores of industries and governments worldwide. The *Wall Street Journal's* Centennial Issue named him among 39 people in the world most likely to change the course of business in the 1990s, and *Car* magazine, the 22nd most powerful person in the global automotive industry. His work focuses on transforming the car, real-estate, electricity, water, semiconductor, and several other sectors of the economy toward advanced resource productivity. His latest books are *Natural Capitalism: Creating the Next Industrial Revolution* (with Paul Hawken and L. Hunter Lovins, 1999, www.natcap.org) and *Small Is Profitable: The Hidden Economic Benefits of Making Electrical Resources the Right Size* (RMI, August 2002).

Ron Perkins has been involved in the design, construction and operation of commercial and light industrial facilities for the past 30 years. He has a BS in industrial arts from Sam Houston State University with a minor in mathematics. He has worked for Todd Shipyards Corporation, Offshore Power Systems, Texas Instruments, Inc., and Compaq Computer Corporation before co-founding Supersymmetry USA. For eight years, ending in July 1990, Ron Perkins held the position

of Facilities Resource Development Manager at Compaq Computer Corporation. He managed a 50-member design team of architects, engineers, contractors and scientists designing over 3,000,000 square feet of state-of-the-art, commercial office and factory space, housing Compaq Computer Corporation's World Headquarters in Houston, Texas. Perkins formed a team, to research and apply energy efficient technologies. As the result of the team's efforts, Compaq's new buildings cost less to build and are 30% more efficient. For the last 12 years, Ron Perkins has served as president of Supersymmetry USA, Inc., a sustainable mechanical design-consulting firm located near Houston, Texas. Working with design teams on diverse projects ranging from guided missile cruisers to office buildings, Perkins brings integrated design methodology, real-time performance measurement, and whole-systems thinking to the design process.

Peter Rumsey, PE, CEM, is founder and Principal of Rumsey Engineers, an HVAC engineering, design, and consulting firm in Oakland, CA. Peter has over 20 years' experience in the building design field. Peter is a graduate of UC Berkeley with a mechanical engineering degree and is a registered mechanical engineer in six states, a certified energy manager, and a member of the Association of Energy Engineers and ASHRAE. Peter joined Supersymmetry USA in 1996 and worked closely under the guidance of Lee Eng Lock, a world leader in efficient design of mechanical systems. Peter ran and owned Supersymmetry's West Coast office and in 2000 renamed it Rumsey Engineers. Peter is an up-and-coming leader in the HVAC design field. Recently he has received a national award from ASHRAE and was named the energy engineer of the year for the San Francisco Bay area. He specializes in the design of efficient mechanical systems for office buildings, public sector buildings, and critical environments such as cleanrooms, data centers, and laboratories. He has worked on numerous sustainable design and green building projects including buildings designed to meet the LEED rating system. Peter is currently a member of the ASHRAE Clean Spaces (Cleanrooms) Technical Committee 9.11. He is playing an important role in redefining how the mechanical systems in buildings are designed and built. Some of his clients include Netscape, Applied Materials, Intel, LoudCloud, the Carnegie Institution, the City of San Francisco, and Lawrence Berkeley National Laboratory.

Dale Sartor, PE, heads the LBNL Building Technologies Applications Team which assists in the transfer of new and underutilized technology through project focused multi-disciplinary teams. Mr. Sartor has an AB in architecture, and a master's degree in business administration. He is a licensed mechanical engineer, and a licensed general building contractor. He has over 25 years of professional experience in energy efficiency and renewable energy applications including 10 years as a principal of an architecture and engineering company, and seven years as the head of LBNL's in-house energy management program.

Jenifer Seal, principal, is a member of Rocky Mountain Institute's Green Development Services team. She holds a master's degree in real estate development from Massachusetts Institute of Technology, a bachelor of architecture and a B.S. in environmental design from Ball State University. Ms. Seal is a consultant on green development and energy-efficient building. She is a senior coauthor of RMI's 525-page book, *Green Development: Integrating Ecology and Real Estate* and *Green Developments* CD-ROM, and has managed and participated in a number of projects such as the Pentagon renovation charrette and the Pittsburgh Nine Mile Run stormwater charrette. Jenifer was also a managing director of RMI's Natural Capitalism Practice, in which capacity she designed business workshops and seminars and played a key role in the Institute's strategic planning.

Joel Swisher, PhD, PE, is a Principal and Team Leader of Energy & Resources Services at Rocky Mountain Institute. Dr. Swisher is a registered professional engineer and holds a Ph.D. in energy and environmental engineering from Stanford University. With 25 years' experience in research and consulting on many aspects of clean energy technology, Dr. Swisher is an internationally recognized expert in the analysis, design, and evaluation of utility energy efficiency, distributed generation and emission reduction programs, and the development and finance of carbon offset projects. He is currently leading RMI's consulting work with the City of San Francisco to develop a sustainable energy plan and implementation strategy. During the first half of 2003, he is teaching a graduate course in greenhouse gas management in the Civil and Environmental Engineering Dept. at Stanford University.

Tom Watanabe, adjunct marketing specialist, received his BS in business administration and computer science from the University of Southern California. As business development director for RMI's Data Center Charrette, Mr. Watanabe led the effort to make business contacts for this unique workshop, which is being convened to brainstorm the radical reduction of energy use by large server rooms and data processing hubs. Previously, as senior account executive for Forrester Research, Inc. of San Francisco, he brought in new business clients for Forrester's eBusiness research and advisory services. Much of Mr. Watanabe's marketing career has been spent in Asia. He was North Asia Sales Manager and Asia Business Development Manager for MOD-TAP/Molex Premise Networks, working in Japan, China, Hong Kong, and the Philippines. For the same company, as Asia Sales Administrator, based in Melbourne, he overhauled company policies to achieve better customer satisfaction. He was Sales Team Leader, based in Japan, for Linc Computer, and held marketing positions with Tandon Corporation and Moore/Businessland/Sears in the United States.

Geoff Wood is a native of Trinidad and Tobago, currently residing in Sidney, BC, Canada. His company, Profile Composites, is a design and development firm specializing in applications of advanced materials and processes. Their portfolio includes clients in transportation, advanced energy, marine, electronics housings, and aerospace. Prior to this Geoff worked for the Oak Ridge National Laboratory in Tennessee. Efforts there focused on researching and managing projects in a) advanced materials for the DOE automotive lightweight materials program and Partnership for A New Generation of Vehicles (PNGV), b) specialty materials applications for military space-based hardware, and c) advanced low-signature ground structures. Geoff has a bachelor's degree in chemical engineering from Cornell, a master's degree in materials science from the University of British Columbia, and is currently pursuing a Ph.D. in mechanical engineering at the University of Victoria.

Participant List

Barry Abramson

Sr. Vice President, Engineering
Servidyne Systems
LLC 1350 Spring St., NW, Suite 200
Atlanta, GA 30309
404-352-2050 X 255
barry.abramson@servidyne.com

Eric Adrian

Jones Lang LaSalle
200 East Randolph
Chicago, IL 60601
eric.adrian@am.joneslanglasalle.com

Bernard Aebischer

Centre for Energy Policy & Economics (CEPE)
ETH Zentrum, WEC
8092 Zurich
+41-1-632 41 95
www.cepe.ethz.ch
bernard.aebischer@cepe.mavt.ethz.ch

Adrian Altenberg

Anstein+Walther AG
Leutschenbachstrasse 45, 8050
Zürich Switzerland
+41 305 91 11
www.amstein-walther.ch
adrian.altenburger@amstein-walther.ch

Dan Baer

VP, Environmental Products
Member of Telcordia Comm. on Thermal Mgmt.
in DCs
Liebert Corporation
1050 Dearborn Drive
Columbus, OH 43085-1544
614-841-5723
www.Liebert.com
www.cha-portal.com/
Dan.Baer@liebert.com

Kevin Best

Director of Sales & Marketing
RealEnergy, Inc.
300 Capitol Mall, Suite 120
Sacramento, CA 95814
916-325-2500 X 109
707-426-3300
kbest@realenergy.com

Clark Bisel

Senior Vice President
Flack + Kurtz Consulting Engineers
343 Sansome St., Suite 450
San Francisco, CA 94104
415-402-2231
Clark.Bisel@sf.fk.com

Kenneth Brill

Executive Director
The Uptime Institute, Inc.
1347 Tano Ridge Rd.
Santa Fe, NM 87506
505-982-0514
505-577-2098 cell
www.upsite.com
tui@upsite.com

Chris Chouteau

Senior Vice President
Pacific Gas & Electric
P.O. Box 770000
San Francisco, CA 94177
415-972-5239
wcc2@pge.com

Tom Coulard

Data Support Associates, Inc.
P.O. Box 220
Brookfield, CT 06804
203.740.4222
516-236-9198
tcoulard@datasupportassociates.com

Dave Coup

Project Manager
New York State Energy Research and
Development Authority
(NYSERDA)
17 Columbia Circle
Albany, NY 12203-6399
518-862-1090 X 3379
jdc@nyserda.org

Tom Croda

Principal Engineer
ClearBlue Technologies
100 First St., Suite 2000
San Francisco, CA 94105
415-869-7100
tom.croda@clearblue.com

Joe Daniels

Engineer, Facilities Group, Power Section
Department of Defense / National Security Agency
301-688-5560
joedaniels711@netscape.net

Tom Ditoro

HDR Architects
8404 Indian Hills Drive
Omaha, NE 68144
402-399-1169
www.hdrinc.com
tditoro@hdrinc.com

Patsy Dugger

Program Manager
Industrial and Agricultural New Construction
Customer Energy Management
Savings By Design Program
Pacific Gas and Electric Company
P.O. Box 770000
San Francisco, CA 94177
415-973-1019
pwd2@pge.com

Grant Duhon

Program Manager Supervisor
Customer Energy Management
Savings By Design Program
Pacific Gas and Electric Company
Mail Code: N6G
P.O. Box 770000
San Francisco, CA 94177
415-972-5695
gxd7@pge.com

Gary Engle

California ISO
151 Blue Ravine
Folsom, CA 95630
916-608-5974
gengle@caiso.com

Wu-chun Feng

Los Alamos National Lab
CCS-1 MS-D451
LANL
Los Alamos, NM 87545
505-665-2730
feng@lanl.gov

Stephen Fok

Senior Program Engineer
Technical Application Services
Customer Energy Management
Savings By Design Program
Pacific Gas and Electric Company
415-973-4735
SKF2@pge.com

Rafael Friedman

Senior Project Manager
Customer Energy Management
Mail Code N6G
Pacific Gas & Electric Company
P.O.Box 770000
San Francisco, CA 94177-0001
415-972-5799
415-972-4799
RAFi@pge.com

John Gage

Chief Scientist
Sun Microsystems
P.O. Box 2065
Burlingame, CA 94011
650-786-3400
Ass't: Tricia Fitzgerald
patricia.fitzgerald@sun.com
www.johngage.com
john.gage@sun.com

Steve Greenberg

President
Thin Client Computing
34522 N. Scottsdale Rd. D8453
Scottsdale, AZ 85282
602-432-8649
www.thinclient.net
steveg@thinclient.net

Steven Greenberg

Senior Energy Officer
RealEnergy, Inc.
5957 Variel Ave.
Woodland Hills, CA 91367
818-610-2300
sgreenberg@realenergy.com

John Griffiths

ARUP
901 Market St., Suite 260
San Francisco, CA 94103
415-946-0298 direct
415-994-0754 cell
415-957-9445 switchboard
415-957-9096 fax
john.griffiths@arup.com
www.arup.com

Peter Gross

CEO
EYP Mission Critical Facilities Inc.
11845 W. Olympic Blvd.
Suite 850 West
Los Angeles, CA 90064
310-914-3442
www.eypae.com
pgross@eypmcf.com

Don Hall

Account Services Supervisor
Pacific Gas and Electric Company
111 Almaden Blvd
San Jose, CA
408-282-7444
DJH5@pge.com

Chris Hipp

Founder and former CTO
RLX Technologies
founding member
Server Blade Trade Association
1590 Stockbridge Ave.
Redwood City, CA 94061-4222
650-455-8881
chris@hipster.com

Ron Hughes

Principal
California Data Center Design Group
14925 Lago Dr.
Rancho Murieta, CA 95683
916-531-7396
rhughes@cdcdg.com

Jerry Hutchinson

Major Account Rep - Data Centers
Account Services
Pacific Gas and Electric Company
111 Almaden Blvd.
San Jose, CA
408-282-7345
JAHo@pge.com

Ernie Jensen

Utility Manager
Telus Communications
12040-107 St.
Edmonton, Alberta, Canada T5G 2S7
780-493-7610
ernie.jensen@telus.com
www.telus.com

Steve Jurvetson

Managing Director
Draper Fisher Jurvetson
400 Seaport Ct., Suite 250
Redwood City, CA 94063
650-599-9000
www.drapervc.com
SJ@dfj.com

Ron Kalich

Director, Facilities
ClearBlue Technologies
100 First St., Suite 2000
San Francisco, CA 94105
415-869-7116
ron.kalich@clearblue.com

Ray Keane

ARUP
901 Market St., Suite 260
San Francisco CA 94103
415-957-9445
415-957-9096 fax
ray.keane@arup.com
www.arup.com

Mukesh Khattar

Energy Director
Oracle Corporation
500 Oracle Parkway, MS LGN-2
Redwood Shores, CA 94065
650-506-6980
mukesh.khattar@oracle.com

Jonathan Koomey

Staff Scientist, End-use Forecasting Group
Lawrence Berkeley Nat. Lab (LBNL)
1 Cyclotron Rd., Building 90-4000
Berkeley, CA 94720
510-486-5974
http://enduse.lbl.gov
JGKoomey@lbl.gov

Henry Lau

PE, Design and Engineering Services
Southern California Edison
6042 N. Irwindale Ave., Ste. B
Irwindale, CA 91702
626-633-7179
lauh@sce.com

Donald Lee

Senior Vice President
Bank of America
#10609
CA5-701-10-07
1455 Market St., 10th Floor
San Francisco, CA 94103
donald.t.lee@bankofamerica.com

Jim Magdych

CIO
Cool Chips, plc.
408-621-6125
jmagdych@coolchips.gi

KC Mares

Director, Data Center Operations
Redundant Networks
9408 Double R Blvd., Suite B
Reno, NV 89521
775-850-4222 X 2237
www.redundant.com
kcmares@redundant.com

Greg Mears

VP, Engineering
Invensys Powerware
8609 Six Forks Rd.
Raleigh, NC 27615
919-878-1003
greg.mears@invensys.com

Ben Mehta

Major Account Rep - Cleanrooms
Account Services
Pacific Gas and Electric Company
111 Almaden Blvd.
San Jose, CA
408-725-2062
DJH5@pge.com

Bruce Nordman

Lawrence Berkeley Nat. Lab (LBNL)
1 Cyclotron Rd.
90-4000
Berkeley, CA 94720-8136
510-486-7089
BNordman@lbl.gov

John Pappas

Principal
Mazzetti and Associates
530 Bush St., 3rd Floor
San Francisco, CA 94108
415-362-3266
www.mazzetti.com
johnp@mazzetti.com

Chandrakant Patel

P.E., Principal Scientist
Internet Systems and Storage Laboratory
Hewlett-Packard Laboratories
1501 Page Mill Rd.
Palo Alto, CA 94304
650-857-7140
650-743-1599 cell
chandrakant_patel@hp.com

Bob Perreault

Caliber Facilities Mgt. Ltd.
1801 Broadview Rd. NW
Calgary, AB Canada T2N 3H6
403-270-3223
calfac@telusplanet.net

Neil Rasmussen

SVP, Chief Technical Officer
American Power Conversion
85 Rangeway Rd., Bldg. 2
North Billerica, MA 01862
www.apcc.com
neil.rasmussen@apcc.com

Brad Roberts

S&C Electric Co.
Power Quality Products Division
5251 West Franklin Dr.
Franklin, WI 53132
414-423-8776 X 109
broberts@sandc.com

Paul Roggensack

Mechanical Engineer
California Energy Commission
1516 9th St., MS 43
Sacramento, CA 95814-5504
916-654-6560
Proggens@energy.state.ca.us

Art Rosenfeld

Commissioner
California Energy Commission
1516 9th St., MS-35
Sacramento, CA 95814
916-654-4930,
Arosenfe@energy.state.ca.us

David Schirmacher

Goldman Sachs
180 Maiden Lane
15th Floor
New York, NY 10038
212-357-7482
david.schirmacher@gs.com

Roland Schoettle

Chief Executive Officer
Optimal Technologies International, Inc.
1513 Spyglass Parkway
Vallejo, CA 94591
707-557-1788
916-718-9274 cell
rolands@otii.com

Steve Schumer

VP, Technology
Cupertino Electric Energy Group (CEI Energy Group)
1132 North Seventh St.
San Jose, CA 95112
408-808-8173
408-966-9550 cell
Steve_Schumer@cei.com

Bob Seese

Manager, Data Center Services
Loudcloud
599 N. Mathilda
Sunnyvale, CA 94085
408-212-4630
www.loudcloud.com
bseese@loudcloud.com

Peter Spark

Principal
Ecsponent
R09022
Flat 8, 37-39
Great Ormond Street
London WC1N3HZ
pspark.ocs2000@london.edu

Mike Steinmann

General Manager, San Francisco
Glumac International
10419 Old Placerville Rd. # 250
San Francisco, CA 94107
415-398-7667
msteinmann@glumac.com
www.Glumac.com

Joe Stolarski

SVP, Director of Engineering and Operations
Jones Lang LaSalle
200 East Randolph
Chicago, IL 60601
312-228-2091
joe.stolarski@am.joneslanglasalle.com

Steve Strauss

President
Glumac International
320 SW Washington, Suite 200
Portland, OR 97204
503-227-5280
www.Glumac.com
sstraus@glumac.com

Harold J. Stewart

Principal
Keen Engineering Co. Ltd.
116 - 930 West 1st St.
North Vancouver, BC
Canada V7P 3N4
604-986-5336
604-983-4034 direct
604-219-8385 cell
harold.stewart@keen.ca

Stephen Torres

VP, Western Region
Fuel Cell Energy, Inc.
225 South Lake Ave., Suite 300
Pasadena, CA 91101
626-432-5410
626-818-4647 cell
storres@fce.com

William L True

FCRE Corporate Engineering
82 Devonshire St., W3B
Boston, MA 02109
617-563-7349
bill.true@fmr.com

William Tschudi

Lawrence Berkeley Nat. Lab (LBNL)
1 Cyclotron Rd. Ms 90R3111
Berkeley, CA 94720
510-486-5988
wftschudi@lbl.gov

Ashlee Vance

IDG News Service
San Francisco Bureau
San Francisco, CA 94107
415-974-7391
Ashlee_Vance@idg.com

James Warren

Atlantic Regional Director
EYP Mission Critical Facilities, Inc.
54 State St., 3rd Floor
Albany, NY 12207
518-337-2013
jwarren@eypmcf.com

Bill Westbrook

Business Development Manager
Power Measurement USA
1099 D St. Suite 208
San Rafael CA 94901
877-476-5872
415-457-9040 direct
415-595-6716 mobile
bill_westbrook@pml.com

Scott Wheeler

Principal
The Engineering Enterprise Consulting Engineers
853 Lincoln Way, Suite 109
Auburn, CA 95603
530-886-8556
tee-sw@pacbell.net
www.engent.com

Richard Williams

Jones Lang LaSalle
200 East Randolph
Chicago, IL 60601
richard.williams@am.joneslanglasalle.com

John Wilson

Special Ass't to Art Rosenfeld
California Energy Commission
1516 9th St.
Sacramento, CA 95814-5512
916-654-5056
jwilson@energy.state.ca.us

Ron Wilson

Office Director
M&A SiliconValley
452 South First St.
San Jose, CA 95113
408-279-2390
RonW@mazzetti.com

Tim Xu

Project Manager, Building Technologies Department
Environmental Energy Technologies Division
Lawrence Berkeley Nat. Lab (LBNL)
One Cyclotron Road
BLDG 90R3111
Berkeley, CA 94720-8134
510-486-7810
ttxu@lbl.gov

note: not all participants submitted bios
speakers/panelists are denoted in red

Barry Abramson, PE, has devoted his professional life to improving the energy efficiency, environmental impact, and economic performance of buildings. As Sr. Vice President of Servidyne Systems, LLC, based in Atlanta, Georgia, he manages an engineering staff that provides technical consulting services for commercial and institutional facilities. Specializing for over 20 years in building energy efficiency, he has consulted for the California Energy Commission, Equity Office Properties, Lend Lease Real Estate, Time-Warner, the Sears Tower, Ritz-Carlton Hotels, and many other government and private-sector organizations throughout the United States. Abramson graduated Phi Beta Kappa with a degree in energy and environmental engineering from Stanford University in 1978. Before joining Servidyne in 1981, he worked for the City of Palo Alto municipal utility, helping develop its energy conservation programs. He is a registered professional engineer in 12 states.

Bernard Aebischer earned a Ph.D. in physics and has worked for more than 20 years as an energy analyst and policy consultant. For over 10 years he has served as a senior scientist with the Swiss Federal Institute of Technology in Zurich and staff member of CEPE—the Centre for Energy Policy and Economics (www.cepe.ethz.ch)—since its foundation in 1999. Past and future energy use and energy efficiency in the tertiary sector and by ICT-equipment are his first research fields. On behalf of the Swiss Federal Office of Energy, he is leading the competence center “Energy and Information Technology.” In January 2003 together with engineers and policy analysts he published the report “Energy- and Eco-Efficiency of Data Centres,” a study commissioned by the Canton of Geneva (15 Jan 2003). www.cepe.ethz.ch/news/news/index.html

Adrian Altenburger is Partner and Head of the HVAC Department at Amstein+Walthert Ltd., an independent engineering company with 265 employees in the field of HVAC and electrical engineering, consultants for security, facility management and building physics, based in Zurich, Switzerland. His specialization includes: HVAC and energy system Concepts in the field of data centers and office buildings project management for M&E. In 2000, Adrian earned a diploma in general building project management from the Swiss Federal Institute of Technology ETH, Department of Architecture, in Zurich, Switzerland.

Dan Baer is Vice President for Environmental Products, specializing in telecom and server solutions for Liebert Corp. Mr. Baer has been a Liebert Associate for over 26 years. He has held several key leadership positions within Liebert, including Vice President of Environmental Products, Vice President Communications Technology Group, Vice President of Product Management, and Director of Engineering. Mr. Baer has a BS in mechanical engineering from Marquette University in Milwaukee, WI. Other degrees and accomplishments include a bachelor of management engineering degree from Rensselaer Polytechnic Institute in Troy, NY, certification as a registered professional engineer in the state of Ohio, and the holder of two U.S. patents.

Kevin D. Best, Principal, is a founding partner of RealEnergy, with over 25 years in real estate and energy development. He has lead several energy efficiency and distributed generation developments for institutional investors throughout their facility portfolios. Mr. Best develops, owns and operates energy generation assets; sells power priced below future retail rates; provides solutions for power reliability, quality, price risk management, environmental stewardship goals or process enhancement. Mr. Best is a well-known authority on distributed generation, energy efficiency and competitive electric issues.

Clark Bisel has over 20 years’ of experience with Flack + Kurtz. As an industry leader in sustainable design, his projects have been recognized by their uniqueness in approach and for their energy performance results. Mr. Bisel is also actively involved in the analysis and implementation state-of-the-art technology such as thermal storage, daylighting, passive solar and central heat pump design, and building systems automation. Mr. Bisel’s completed design projects exceed 50 million square feet of space in various building types and systems. These include government, mixed-use, high-rise commercial developments, data centers, museums, and hospitality facilities.

Kenneth G. Brill is a management consultant, executive director of The Uptime Institute, and president of Computersite Engineering. He is the originator of the Business Case for Site Infrastructure Reliability. Mr. Brill’s effective approach focuses on the managerial, engineering and strategic differences between systems that “work” and complex infrastructures that never fail, or fail transparently without affecting users. His expert-systems reliability rules and analytical processes systematically identify infrastructure vulnerabilities, which are normally discovered only in the aftermath of a downtime disaster. Mr. Brill is the creator of the site uptime reliability matrix, the articulator of needed site infrastructure capabilities, including permanently imbedding intelligence within equipment, the automatic execution of manual processes, and virtual infrastructure training simulators. He holds a patent on dual power technology, and has received several industry awards.

David Coup has served as a project manager for the New York State Energy and Research and Development Authority (NYSERDA) since 1999. There, he develops public benefit programs that help electric service customers and providers manage peak electricity demand. These programs lend stability to the electric grid, and its market pricing, during periods of high demand. Also, Mr. Coup develops programs to demonstrate the next generation of energy efficient, end-use technologies, and markets them to various stakeholders in New York state, including those involved with data centers. Prior to NYSERDA, Mr. Coup spent 11 years leading programs to increase productivity, enhance financial management processes, and improve environmental practices within several divisions of the General Electric Company. Mr. Coup holds a bachelor’s degree in industrial and systems engineering from Ohio State University, and a master’s degree in environmental management and policy from Rensselaer Polytechnic Institute.

Thomas (Tom) G. Croda, Principal Engineer for Navisite, has more than 34 years' experience in telecommunications and is responsible for standards related to DC and AC power and other common systems areas. He was Principal Engineer for Sprint Long Distance for 17 years, responsible for standards related to DC and AC power and other common systems areas. He majored in electronic engineering at California State Polytechnic College. He is the recipient of numerous commendations for distinctive innovations and the developer of various products contributing to connectorization, improved grounding techniques, and DC power plant design. He was instrumental in establishing equipment designs for solar power plant components. The past Vice Chairman of Technical Sub-committees T1E1 and T1Y1, he is the present Convener of Working Group T1E1.5 Telecom Power. He is a member of the IEEE Power Engineering Society Committee on Stationary Batteries and an advisory Board Member of the Protection Engineers Group.

Grant Duhon is Pacific Gas & Electric's Supervising Program Manager for nonresidential new construction programs. Since 1993, Mr. Duhon has served the customers of Pacific Gas & Electric by promoting conservation, energy efficiency, and integrated building and systems design. He feels fortunate to have been involved with projects that have advanced the science of integrated design. Grant was part of the statewide team that developed the Savings By Design and Energy Design Resources programs. Grant has worked in the new construction industry since 1976, entering the industry through the trades. After receiving his undergraduate degree in Nevada, he continued his involvement in the industry as an HVAC sales engineer and consultant, working for the Trane Company and Lennox Industries, among others. Grant's first contact with the value of energy efficiency in commercial construction came in 1980 with the design and installation of a major project which paid for itself in energy savings in less than 90 days. Excited by this success, Grant has since promoted energy and resource efficiency as a central consideration in design and construction.

Thomas Ditoro is a registered electrical engineer in the state of Nebraska, and is a Project Electrical Engineer for HDR, Inc. Mr. Ditoro has extensive mission-critical facility design experience. Representative projects have included fuel-cell-powered data centers, colocation data centers, nanotechnology facilities, and healthcare facilities. Before joining HDR, Inc., he was Facilities Engineering Manager for Southwestern Bell's data centers in Dallas and Houston, Texas. He also served as the Chief Technology Officer of 7X24 Facilities, a start-up colocation data center company. He currently is a member of the Gerson-Lehrman Council of Advisors for fuel cell and distributed generation technologies.

Wu-chun Feng received a BS degree in computer engineering and a BS (honors) degree in music from Penn State University in 1988; an MS degree in computer engineering from Penn State University in 1990; and a Ph.D. degree in computer science from the University of Illinois at Urbana-Champaign in 1996. Dr. Feng is currently a technical staff member and team leader of RADIANT (Research & Development in Advanced Network Technology) at Los Alamos National Laboratory and an adjunct assistant professor at Ohio State University. He is a fellow of the Los Alamos Computer Science Institute and the founder and director of the Advanced Summer Curriculum for Emerging Network Technologies (ASCENT). Before joining LANL in 1998, Dr. Feng had previous professional stints at Purdue University, the University of Illinois, NASA Ames Research Center, and IBM's T.J. Watson Research Center.

John Gage is the Chief Researcher and Director of the Science Office, for Sun Microsystems, Inc. He is responsible for Sun's relationships with world scientific and technical organizations, for international public policy and governmental relations in the areas of scientific and technical policy, and for alliances with the world's leading research institutions. Gage attended the University of California, Berkeley, the Harvard Kennedy School of Government, and the Harvard Graduate School of Business. He did doctoral work in mathematics and economics at the University of California, Berkeley. He is a member of the Mathematical Association of America, the Association for Computing Machinery (ACM), the Institute of Electrical and Electronics Engineers (IEEE), and the Board of Trustees of the Internet Society (ISOC).

Steve Greenberg is the founder and President of Thin Client Computing in Scottsdale, Arizona. He is the author of a recent study "Power to the People: Comparing Power Usage for PCs and Thin Clients in an Office Network Environment." A leading expert in server-based computing solutions he has designed mission-critical solutions for various *Fortune* 500 companies.

Steven A. Greenberg is Chief Operating and Senior Energy Officer for RealEnergy of which he is also a founder. As COO/SEO he is responsible for the company's design, construction, plant operations, and government and regulatory affairs. He has 18 years' of experience in the electric energy and utility industry. He started his career at PG&E, where he held a variety of positions in power plant operations, business development, government relations, and power contracts. Prior to founding RealEnergy, Mr. Greenberg was a cofounder and managing director of Intergy, LLC, a predecessor of RealEnergy. Other industry experience includes employment as a senior consultant at Henwood Energy Services where he had primary responsibility for the firm's qualifying facility and new project generation development practice, as well as substantial tenure at a large investor-owned utility where he held positions in power plant operations, project management, business management, government relations, and power contracts. Mr. Greenberg has extensive experience dealing with government, utility, and industry trade groups within the energy sector and can often be found speaking before state and national organizations regarding the restructuring of the utility industry and the advancement of distributed generation. He is an Executive Board member of the U.S. Combined Heat and Power Association and has served on the CADER Executive and Steering Committees and the Executive Board of the Distributed Power Coalition of America. Mr. Greenberg has a BS in Business Administration from California Polytechnic Institute.

Peter Gross has over 24 years of experience in the engineering, design and operational support of high-reliability infrastructure systems. He currently serves as Chief Executive Officer and Chief Technology Officer of EYP Mission Critical Facilities, Inc., one of the largest engineering and consulting firms focused on the design and construction of data centers, trading floors, communication, and broadcasting facilities. Since joining the firm in 1997, he has been actively involved in the development and the rapid growth of the Mission Critical Engineering practice. His present responsibilities include strategic planning, technical oversight, and business development. In the past 10 years, Mr. Gross has managed various projects for numerous *Fortune* 500 Companies such as AOL, Exodus Communications, AT&T, American Airlines, Bankers Trust, IBM, Wells Fargo Bank, Intel, Pacific Bell, Shell, Merrill Lynch, Charles Schwab, Fidelity Investments, IRS, Fox Television, WorldCom MCI, Southwestern Bell, Citigroup/Salomon Smith Barney, JP Morgan Chase, Bank One, and Sprint.

Chris Hipp, cofounded RLX Technologies with Messrs. Harkey and Cracken in December 1999. Mr. Hipp invented the concept of the “serverblade” and the patented architecture of the original RLX System 324. Mr. Hipp served as RLX’s Chief Technology Officer through the Spring of 2001. More recently, Mr. Hipp served as a market and technology evangelist for RLX. In May 2001, RLX Technologies beat everyone to the punch by bringing the first ultra-dense bladed server solution to market. By creating both ServerBlades and the software required to manage them, RLX has become the leader in the emerging “blade” computing market.

Ron Hughes has been involved in the design, operation, construction, testing, and commissioning of data centers for over 20 years. In the last five years alone, as owner/principal of California Data Center Design Group, Mr. Hughes has supervised the design of over 1,200,000 square feet of state-of-the-art data centers in six different countries. The data centers Mr. Hughes has designed include corporate, governmental, financial-grade, and collocation data centers. Prior to his full-time focus on data center design and construction supervision, Mr. Hughes served as the Manager of Facilities, Engineering and Technical Planning for a 160,000-s.f. data center in Sacramento, California. A frequent speaker at national conferences, Mr. Hughes recently presented his views on building data centers in foreign countries at the 2002 7X24 Spring conference in Orlando, Florida. He is also the author of numerous technical articles including: “Designing Data Centers for the 21st Century, Reliability Studies, What Should a Client Expect?,” “The Top Ten Ways to Keep Your Data Center Online,” “Protecting Your Data Center or Critical Facility from Terrorist Attack,” “The 100 Watts Per Square Foot Data Center, Reality or Myth?,” and “Data Center Disasters and How to Avoid Them.”

Steve Jurvetson is a Managing Director of Draper Fisher Jurvetson. He was the founding VC investor in Hotmail (MSFT), Interwoven (IWOV), and Kana (KANA). He also led the firm’s investments in Tradex (acquired by Ariba for \$6B) and Cyrus (acquired by Ciena for \$2B), and most recently, in pioneering companies in nanotechnology and molecular electronics. Mr. Jurvetson was an R&D engineer at Hewlett-Packard, where seven of his communications chip designs were fabricated. His prior technical experience also includes programming, materials science research (TEM atomic imaging of GaAs), and computer design at HP’s PC Division, the Center for Materials Research, and Mostek. At Stanford University, he finished his BSEE in 2.5 years and graduated first in his class, as the Henry Ford Scholar. Mr. Jurvetson also holds an MS in electrical engineering from Stanford. He received his MBA from the Stanford Business School, where he was an Arjay Miller Scholar. Mr. Jurvetson also serves on the Merrill Lynch Technical Advisory Board and is Co-Chair of the NanoBusiness Alliance. He was recently honored as “The Valley’s Sharpest VC” on the cover of *Business 2.0* and chosen by the *San Francisco Chronicle* and *San Francisco Examiner* as one of “the ten people expected to have the greatest impact on the Bay Area in the early part of the 21st century.”

Ron Kalich has over 15 years’ experience in the information systems, communications, power, and—now converged—data center industries. Ron has worked for Ameritech, Pacific Gas & Electric, and most recently several data center operating companies in finance, regulatory, engineering, and operating positions. He’s been responsible for the design, construction, and operation of more than 30 data centers in the United States and abroad. Ron currently works for Navisite as Director-Facilities Engineering, overseeing the operations of 16 diverse data centers in the United States.

Jonathan Koomey is a Staff Scientist and Group Leader at Lawrence Berkeley National Laboratory, in Berkeley, California. He holds MS and Ph.D. degrees from the Energy and Resources Group at the University of California at Berkeley, and an A.B. in History of Science from Harvard University. He is the author or coauthor of seven books and more than one hundred and thirty articles and reports. His latest book is *Turning Numbers into Knowledge: Mastering the Art of Problem Solving* (www.numbersintoknowledge.com). Koomey serves on the Editorial Board of the journal *Contemporary Economic Policy*, and has appeared on Nova/Frontline, BBC radio, CNBC, All Things Considered, Marketplace, Tech Nation, On the Media, the California Report, KPIX TV (SF), CNET radio, and KQED radio. He has been quoted in the *Wall Street Journal*, *Barron’s*, the *Washington Post*, *Science*, *Science News*, *American Scientist*, Dow Jones News Wires, *USA Today*, *San Francisco Chronicle*, the *Oakland Tribune*, *Interactive Week*, *Business 2.0*, Salon.com, and *Network Magazine*.

Henry Lau received his Ph.D. in mechanical engineering from Duke University in 1973. He moved to California in 1994 to practice in the design of energy-efficient systems for buildings. He is a licensed professional engineer in California. He was in private practice for 19 years. During those years, he was a consultant to the California Energy Commission in the development of the California Energy Standards, known as the Title-24 Building Energy Standards. In 1992, Dr. Lau joined the Energy Efficiency Division of Southern California Edison. He spent his entire time there working on demand side management (DSM). His work included assisting SCE customers to reduce their utility bill by utilizing the latest energy technologies in building energy systems. In 1997, Dr. Lau went to China to represent SCE and spent two weeks lecturing on DSM to the Chinese Ministry of Electric Power. Currently, he is field-demonstrating emerging technologies of energy-efficient buildings and building energy systems.

Jim Magdych is CIO of Cool Chips. Cool Chips, plc. has devised "Cool Chips," which use electrons to carry heat from one side of a vacuum diode to the other. The system, which is currently under development, contains no moving parts or motors and can be miniaturized for use in micro-electronic applications. In almost every electronics cooling application, an active cooling system powered by Cool Chips Technology will be far superior to the more conventional passive solutions. (See www.coolchips.com.)

K.C. Mares is Director of Operations for Redundant Networks, a second-generation data center company. He is responsible for all security and physical infrastructure services, assets, and operations. As a member of the Executive Team, he develops strategic company direction, products, and services. Mr. Mares was previously the Director of Energy and Special Projects at Exodus, a Cable & Wireless service. While at Cable & Wireless, Mr. Mares reported key company metrics to the Cable & Wireless Executive Team, negotiated large customer contracts, and led outsourcing arrangements for utility and accounting services while managing all utility operations.

Bruce Nordman has been with LBNL since 1986, and has focused since 1995 on energy implications of IT equipment. In addition to estimating equipment operating patterns, annual consumption, and savings potentials, he has studied the energy flows embodied in office paper use and more recently how improved user interfaces can save electricity.

John Pappas was educated at California Polytechnic State University in environment engineering with an emphasis in HVAC and solar design, and is a registered engineer in 47 states. Presently, he is a Principal of Mazzetti & Associates in San Francisco, a 100-person engineering design and consulting firm with offices in Minneapolis and California, and practicing worldwide. He has had the pleasure of working in both the construction and consulting engineering fields for the past 23 years. During his career, John has developed and implemented substantial and complex work for a broad range of clients in the corporate, mission-critical, health care, laboratory, and higher education industries, with projects totaling over \$300M in construction. For the past 10 years, John has focused his efforts on the conception, business case justification, design, implementation, testing and operation of mission-critical facilities. He has served such clients as MasterCard, Microsoft, EDS, Goldman Sachs, Visa International, Equinix, Washington Mutual, Intuit, Kaiser Foundation Hospitals, Silicon Graphics, and 3Com. John has pioneered the development of cooling systems for high-density data centers exceeding 200 watts/sf. He is a regular participant and speaker at the 7X24 conferences. He is also a member of the 7X24 Exchange Server Work Group, exploring future trends in server technologies and alternative cooling systems.

Chandrakant Patel is a principal scientist at Hewlett-Packard Laboratories responsible for strategically engaging in thermo-mechanical research for future microprocessors, servers, and data centers at Hewlett-Packard Laboratories. His current interest is research in data center cooling energy consumption at a global level through the HP Labs Smart Cooling Proposition. The Smart Cooling vision is to provision cooling commensurate with the heat loads in a data center, and to provision computing, and thus the heat loads, based on the available cooling resources. The vision is to realize a savings of 50% in cooling energy costs in the global data center network of tomorrow through combination of mechanical thermo-fluids engineering and computer science.

Bob Perreault, of Caliber Facilities Mgt, Ltd is a registered Architect specializing in the design and construction of data centers, server rooms, and communications infrastructure. Rooms range from simple to N+1 complexity. In addition, he provides server reviews and recommendations for established sites (server rooms ranging in size from a few hundred sq. ft., to several thousand sq. ft.). Bob's past experience includes facilities management, data center management, and building construction. Clients include the private sector, utilities, government, and the oil and gas industry.

Neil Rasmussen is Senior VP, CTO, and founder of American Power Conversion Corp. APC is a \$1.5B company focused on power, cooling, and rack infrastructure for critical networks, and is the world largest supplier of Uninterruptible Power Systems. Neil directs the R&D effort at APC and the next-generation data-center design initiative at APC. Neil received his BS and MS degrees from MIT, with a specialty in power electronics. Before starting APC in 1981, he worked at MIT's Lincoln Laboratory on solar-electric power systems and high speed fly-wheel storage systems.

Bradford (Brad) Roberts, Director of Marketing, S&C Electric Company Power Quality Products Division, has over 30 years' experience in the design and operation of critical power systems, ranging from single-phase UPS systems to medium-voltage applications. He began his engineering work as a systems reliability engineer in the Apollo Lunar Module Program at Cape Kennedy. He held senior management positions with two of the major UPS manufacturers during his career. Brad is a member of IEEE and has published over 30 technical journal articles on critical power system design. Brad is a registered professional engineer and has a BSEE (Bachelor of Science in Electrical Engineering) degree from the University of Florida. He is Vice Chairman of the IEEE Power Engineering Society's Emerging Technologies Committee and a member of the Board of Directors for the Electricity Storage Association.

Paul Roggensack is a mechanical engineer at present at the California Energy Commission working on the Public Interest Energy Research (PIER) program. The PIER program funds research, development, and demonstration projects to promote environmentally-safe, affordable, and reliable energy products and services for the State of California. Among the PIER projects he manages is “Energy Efficient Data Centers” (with Lawrence Berkeley National Laboratories) to benchmark energy end use and develop a roadmap to guide future research to enhance energy efficiency at data centers. Other projects include developing distributed generation at oil fields, acoustic stimulation at aluminum foundries, advanced distillation at oil refineries, and a roadmap for energy efficiency at water and wastewater utilities. Prior to joining the Energy Commission, he was a water resources control engineer at the State Water Resources Control Board working on treatment and public works projects to address water pollution. He has also worked as a technician at National Semiconductor in Santa Clara, CA and Acurex Corporation in Mountain View, CA. He has a BS in chemical engineering from San Jose State University.

Joe Stolarski is Senior Vice President Head of Engineering & Operations with Jones Lang LaSalle Americas, Inc., the world’s leading real estate services and investment management firm (www.joneslang-lasalle.com).

Stephen Torres is the Vice-President of the Western Region for FuelCell Energy. FuelCell Energy is a world leader in the development of fuel cell generators for stationary applications—a power generation technology that is among the cleanest and most efficient available for the 21st century. In this role, Mr. Torres is responsible for all FuelCell Energy’s activities in the western United States. He focuses on developing strategic alliances with large energy service providers that want to play a significant role in commercializing FuelCell Energy’s fuel cell power plants and creating a regulatory environment conducive to wide scale deployment of FuelCell Energy power plants. Prior to joining FuelCell Energy, Mr. Torres spent three years at Capstone Turbine Corporation, a leading developer of low-emission, compact power generating MicroTurbine systems, as its Director of Distribution Channels. Earlier in his career, Mr. Torres worked for Deloitte Consulting, a leading worldwide management consulting firm, advising primarily manufacturing clients on supply chain and sales management issues. He has also held marketing and sales positions with Procter and Gamble and General Electric. Mr. Torres received a mechanical engineering degree from the University of Washington and an MBA from the Anderson School of Management at UCLA.

Bill True is in the engineering group at Fidelity Corp. Real Estate working on support base-building and mission-critical operations, capital projects, emergency response, troubleshooting, corporate standards, and new initiatives. He previously managed engineering operations for Boston for five years. He serves as President of the Boston Chapter of 7X24 Exchange. Bill has expert-level experience in HVAC, BAS, electrical, fire protection, and plumbing.

William Tschudi is with Lawrence Berkeley National Laboratory’s Environmental Energy Technologies Division, Energy Analysis Department. His most recent projects include Energy Efficiencies in Laboratory Type Facilities—Clean Rooms (<http://ateam.lbl.gov/clean-room/>) and Benchmarking Energy Use in Cleanrooms. Mr. Tschudi is a licensed professional engineer with 20 years in the power industry and 10 years’ high-tech facilities design. He is a member of ASME and ASHRAE.

Ron Wilson has been the director of the San Jose office of Mazzetti & Associates since 1998. He is a Principal in the firm which totals 100 people and includes offices in San Francisco, San Jose, Sacramento, Los Angeles, and Minneapolis. Ron has 22 years of experience in electrical engineering with emphasis on design, construction, and management of complex projects. With a degree in construction management from California Polytechnic State University in San Luis Obispo he brings a unique and practical approach to design and construction. Through the course of his career, Ron has been involved in all phases of mission-critical, higher education, commercial, industrial, institutional, military, and health care construction. In the past eight years, Ron has been responsible for programming and design of data center facilities including: the Western Operations Center Facility for AOL/Netscape; Replacement and upgrade of the electrical plant at EDS’s Service Management Center in Rancho Cordova, California; the new telephone switch and colocation facility for MCI-WorldCom in Milpitas; the new Campus telecommunications center for California State University at Bakersfield; and programming of the 2.2 million-square-foot, mission-critical campus for US DataPort in North San Jose. In addition to project design and construction, Ron is regularly engaged to evaluate facilities as part of due diligence efforts and to investigate failures in critical facilities. Clients served in this capacity include Sun Microsystems, Kaiser Foundation Hospitals, MCI-WorldCom, Washington Mutual, and Amdocs.



Sponsors

General

Pacific Gas & Electric
www.pge.com

California Energy Commission
www.energy.ca.gov

New York State Energy Research and Development
Authority (NYSERDA)
www.nyserdera.org

ADEME
(Agence de l'Environnement et de la
Maîtrise de l'Energie)
Vanbonne, France
www.ademe.fr

Southern California Edison Company
www.edison.com

BC Hydro Power Smart
www.bchydro.com

The Canton of Geneva (Switzerland)

- Département de l'intérieur, de l'agriculture et de
l'environnement (DIAE)
www.geneve.ch/diae/welcome.asp
- Service cantonal de l'énergie (ScanE)
www.geneve.ch/scane/home/welcome.asp

Centre for Energy Policy & Economics (CEPE)
Zurich, Switzerland
www.cepe.ethz.ch

California Independent System Operators (ISO)
www.caiso.com

In Kind

Lawrence Berkeley National Lab
(LBNL)

John Gage, Sun Microsystems
www.johngage.com

Oregon Office of Energy
www.state.or.us

The Natural Step
www.naturalstep.org

Natural Resources Defense Council (NRDC)
www.nrdc.org

The Uptime Institute, Inc.
www.upsite.com

Industry/Participant Support

American Power Conversion
www.apcc.com

RealEnergy, Inc.
www.realenergy.com

S&C Electric Company
Power Quality Products Division
www.sandc.com

Wu-chun Feng
Los Alamos National Lab

EYP Mission Critical Facilities, Inc.
www.eypae.com

Glumac International
www.Glumac.com

Jones Lang LaSalle Americas, Inc.
www.joneslanglasalle.com

The Engineering Enterprise
www.engent.com



**Appendix XV - 8-22-02 Workshop Agenda, presentation, and
meeting summary**



ERNEST ORLANDO LAWRENCE
BERKELEY NATIONAL LABORATORY

8-22-02 AGENDA ENERGY EFFICIENT DATA CENTERS RESEARCH ROADMAP

☼ INTRODUCTIONS

☼ OVERVIEW OF LBNL DATA CENTER AND RELATED ACTIVITIES

- ☼ IMPACT OF INTERNET ECONOMY**
- ☼ CA DATA CENTER LOAD CHARACTERIZATION**
- ☼ CA DATA CENTER BENCHMARKING/CASE STUDIES (5)**
- ☼ NY DATA CENTER CASE STUDY**
- ☼ IDENTIFICATION OF EFFICIENCY OPPORTUNITY**
- ☼ RESEARCH ROADMAP**

☼ BRAINSTORMING - BARRIERS

☼ BRAINSTORMING – RESEARCH NEEDS

☼ OUTLINE OF DATA CENTER RESEARCH ROADMAP

☼ COLLABRATION OPPORTUNITIES

☼ NEXT STEPS

Barriers To Improved Energy Efficiency in Data Centers:

1. Demands for high reliability and redundancy requires excess UPS capacity causing high losses (inefficiency).
2. Owner wants flexible design, therefore designs for worst possible future load yielding inefficient oversizing of electrical and mechanical systems.
3. Increasing computational output (and even improving efficiency of computation per watt) while increasing energy use per square foot.
4. Fast track scheduling
5. First cost drivers
 - ◆ Connection to UPS - size (and cost) not fully recognized
6. Lack of focus on energy (other much more important drivers)
7. Manufacturers not motivated to produce efficient products (flops per watt)
 - ◆ Energy consumption/efficiency not a market driver
8. Industry does not recognize benefits of DC power distribution – AC considered “given” by server industry (DC accepted in telecommunication industry)
9. Specifications are prescriptive vs. performance
10. Power density capacity is considered a marketing benefit (Bigger is better even if not needed)
11. Likewise, other costly energy specs such as AC capacity, reliability, temperature and humidity control, are used by customers and vendors in selecting and marketing space
12. Poor layout and configuration of computing equipment creates hot spots driving inefficient AC
13. Lack of benchmarking
 - ◆ All systems
14. Information on results of demonstrations and best practice not well captured and distributed
15. Lack of dynamic IT power management
 - ◆ Machine/component level
 - ◆ Room level
 - ◆ Geographical level
16. Case/rack design not optimized for thermal efficiency
17. Lack of standards and communication between hardware and rack manufacturers
18. Over capacity in industry results in inefficient part load operations
19. Focus on high density (why not spread out load?)
20. Co-location “model” built fast and cheap
21. Hard to track (cross cuts SIC codes) - hard to attract interest by public goods funders
22. Lack of metering to separate data center from other functions
23. Lack of accurate information - hard to provide guidance

24. Information highly scattered
25. Technical and organizational support not readily available
26. Difficulty of reaching key decision makers
 - ◆ Designers
 - ◆ Owners
 - ◆ Operators
27. Contractual barriers
28. Customers not charged for capacity
29. Problems with existing benchmarking/testing protocols

R&D Needs

1. Evaluate blade server efficiencies
2. Evaluate other emerging technology for computer hardware
 - ◆ Servers
 - ◆ Routers
 - ◆ Switches
 - ◆ Data storage
3. Look at individual components (above) but also the interconnection/interaction
4. Standardize test for computational energy efficiency
 - ◆ Review existing benchmarking protocols (some problems exist)
 - ◆ Look at “system” (IT hardware and building)
5. Improve transparency of calculations of growth projections
6. Increase benchmarking activity
7. Consider efficiency relative to emerging direct cooling technology
8. Develop other hardware improvements
 - ◆ Power supplies
 - ◆ DC distribution
 - ◆ UPS (and alternatives, e.g. base load CHP)
9. Synergistic options for reliability and energy efficiency
10. Optimizing reliability requirements (how much is enough)
11. Options to on-site reliability (how define reliability)
12. Independent evaluation of reliability levels
13. Better integration of controls
 - ◆ Reduce simultaneous heat/cool, and humidify/dehumidify
 - ◆ Simplify and reduce cost
14. Integrate energy and reliability simulation models (including CFD)
15. Human factors
 - ◆ Overcome barriers
 - ◆ Decision process (and how to influence)
 - ◆ Motivation
 - ◆ Certification/training programs for operators
 - ◆ Improved procedural directions

16. Assess industry trends relative to energy efficiency and look for opportunities to influence.
17. Independent testing and publication of efficiency
18. Improve coordination between private and public organizations/activities
19. Radical change in configuration of data centers
 - ◆ Track trends
 - ◆ Web hosting – centralization (what can it mean) vs. decentralization (for security)
20. Evaluate time scales and integration of various “events”
21. More fully develop “charrette” concept

Related Activity

- iTherm
- CEETHERM
- Uptime institute
- 7X24 Exchange
- HP
- AFCOM (management professional)
- EPRI – Consortium for Electric Infrastructure to Support a Digital Society (CEIDS)
- SVMG
- Utilities
- Telecom Association – USTA
- NYSERDA

Next Steps

1. Circulate meeting notes and attendee list
2. Advisory panel brainstorming on data centers of the future starting with a clean slate
3. Consider collaboration with DOE and the States through STAC
4. Issue draft roadmap and solicit input and priorities
 - ◆ Identify sub groups/topics for focused input
5. Coordination with related groups and activities
6. Encourage RMI to host conference call with potential funders to discuss Charrette
7. Energy efficient data center web site


Energy Efficiency Data Centers



Bill Tschudi
WFTschudi@lbl.gov

Dale Sartor
DASartor@lbl.gov

August 22, 2002



Data Center and Related Activities

◆ Prior Work

- Office Equipment Power Consumption
- Impact of Internet Economy
- Case Study of Data Centers

◆ California Energy Commission


- Data Center Load Characterization for California
- Energy Benchmarking and Case Studies (5)
- Research Roadmap for Data Centers
- Identify Efficiency Improvement Opportunities

◆ NYSERDA

- Case Study of NY Data Center
- Workshop for NY Firms




Electricity used by the Internet




Source: LBNL-44698

Latest estimate of office equipment electricity use by sector in 1999



<http://enduse.lbl.gov/Projects/InfoTech.html>

Comparison of Estimates of Energy Use for Commercial Office Equipment in 1999



Upper bound estimate to electricity used by data centers in the U.S

	Units	1999	2000	2005
Floor area at end of year	Msf	7	20	50
Data center total power density	W/sf	100	100	100
Data center total power	GW	0.7	2.0	5.0
Data center electricity use	TWh	6	18	44
<i>Data centers as % of total use</i>		0.2%	0.5%	1.2%
<i>Average annual contribution to demand growth</i>			0.3%	0.14%
Total U.S. electricity use	TWh	3288	3360	3647

(1) Floor area taken from Juarez, Richard A. 2000. *Virtual Bricks II: Virtual Econ 101 Update: A Comprehensive Guide for Understanding eCommerce Infrastructure Evolution and Convergence*. Boston, MA: FleetBoston Robertson Stephens Inc. May.


(2) Power density is an upper bound based on recent discussions with experts.

(3) Electricity use calculated assuming 8760 hours per year operation, flat load curve.

(4) Total U.S. electricity use taken from EIA's Annual Energy Outlook 2000.


Contact: JGKooomey@lbl.gov, 510/486-5974.

Something unusual is going on



Source: Joe Romm, Center for Energy and Climate Solutions

Electricity Flows in Data Centers



UPS = Uninterruptible Power Supply

PDU = Power Distribution Unit;

Why a Data Center Research Roadmap?

- ◆ California Energy Commission wants to Quantify the Market in California
- ◆ Data Centers are Important to California Economy
- ◆ California wants the “Industry” to Identify Public Goods Research Needs and Priority
- ◆ A Roadmap Will Aid Coordination of Research and Deployment
- ◆ A Roadmap May Stimulate Other Industry Action
 - Uptime Institute
 - ASHRAE
 - iTherm
 - CEETHERM

Roadmap Outline

- ◆ Market Issues
- ◆ Benchmarking
- ◆ Tools
- ◆ Building Systems
 - HVAC
 - Controls
 - Lighting/Other
- ◆ Interface with IT Equipment
- ◆ Efficient Heat Removal in Electronics Equipment

Roadmap Content, con't

- ◆ Deployment
- ◆ Organizations
- ◆ Barriers

Energy End Use in Data Centers




Utility Study Findings



Barrier Examples:

- ◆ Internet Hosting Facilities Demand More Power Capability (And Cooling Capacity) Than is Needed.
- ◆ Conventional Data Center Cooling although Inefficient, is the Industry Standard
- ◆ Design of Processors, Computer Cases, Electronics Racks, and Building Systems are not Integrated

Prioritize the Research Needs



Data Center websites




<http://Datacenters.lbl.gov>

7 X 24 Exchange Energy Efficient Data Centers



Bill Tschudi
October 16, 2002
WFTschudi@lbl.gov



Presentation Outline

- ◆ LBNL Data Center Activities
- ◆ Selected Findings
- ◆ Roadmap Ideas
- ◆ Request Your Input

Data Center Related Activities

◆ Prior Work

- Office Equipment Power Consumption
- Impact of Internet Economy
- Case Study of Data Centers

◆ California Energy Commission

- Data Center Load Characterization for California
- Energy Benchmarking and Case Studies (5)
- Research Roadmap for Data Centers
- Identify Efficiency Improvement Opportunities

◆ NYSERDA

- Case Study of NY Data Center
- Workshop for NY Firms




LBNL Operates NERSC



National Energy Research Scientific Computing Center

Jon Koomey Research into Electricity used by the Internet



Upper bound estimate to electricity used by data centers in the U.S

	Units	1999	2000	2005
Floor area at end of year	Msf	7	20	50
Data center total power density	W/sf	100	100	100
Data center total power	GW	0.7	2.0	5.0
Data center electricity use	TWh	6	18	44
<i>Data centers as % of total use</i>		0.2%	0.5%	1.2%
<i>Average annual contribution to demand growth</i>			0.3%	0.14%
Total U.S. electricity use	TWh	3288	3360	3647

(1) Floor area taken from Juarez, Richard A. 2000. *Virtual Bricks II: Virtual Econ 101 Update: A Comprehensive Guide for Understanding eCommerce Infrastructure Evolution and Convergence*. Boston, MA: FleetBoston Robertson Stephens Inc. May.


(2) Power density is an upper bound based on recent discussions with experts.

(3) Electricity use calculated assuming 8760 hours per year operation, flat load curve.

(4) Total U.S. electricity use taken from EIA's Annual Energy Outlook 2000.

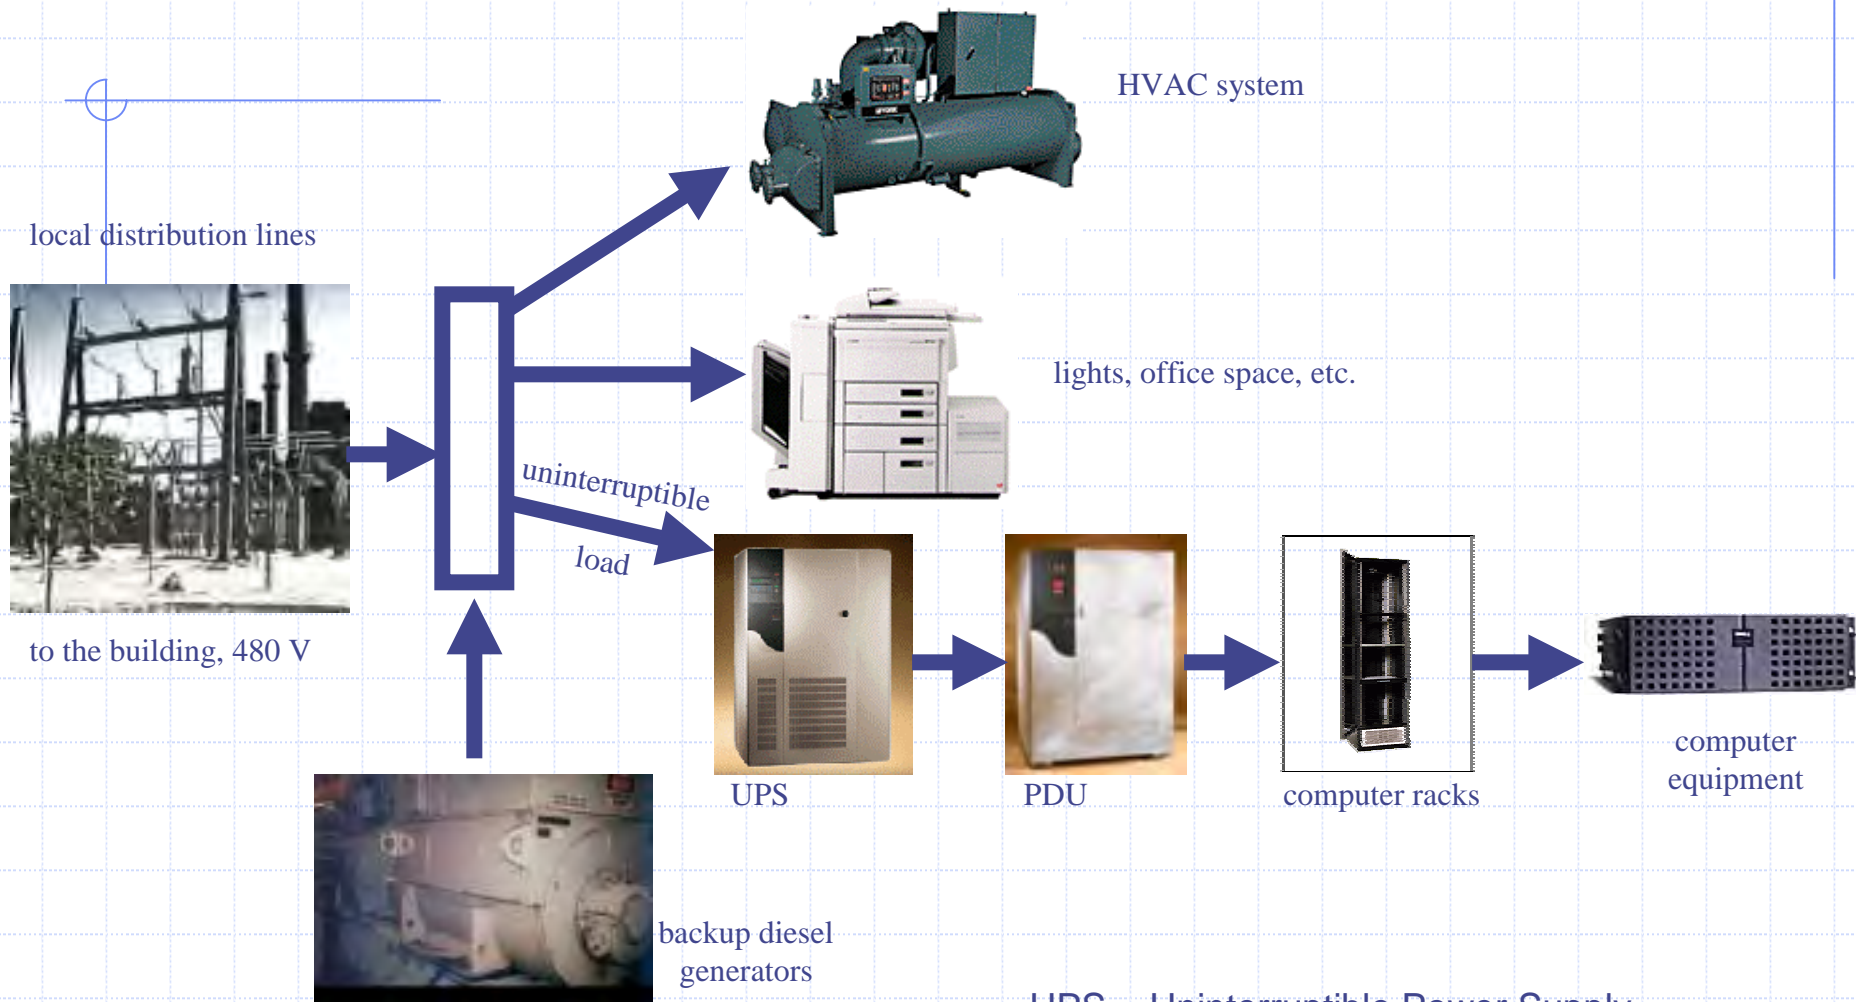
Contact: JGKooomey@lbl.gov, 510/486-5974.

Something unusual is going on



Source: Joe Romm, Center for Energy and Climate Solutions




Electricity Flow in Data Centers



UPS = Uninterruptible Power Supply


PDU = Power Distribution Unit;

Energy End Use From Prior Studies



Source: PG&E Study by Rumsey Engineers


Chart 1. Whole Building Power Consumption



Source: PG&E study performed by Rumsey Engineers


Current Benchmark Results

Energy Balance - Data Center - Facility A. Data Center 1




Current Benchmark Results

Energy Balance - Data Center - Facility A. Data Center 2



Current Benchmark Results

Energy Balance - Data Center - Facility B




Load as a percent of Whole Building

Facility	Whole Building Power Consumption	Data Center Server Load		HVAC Power Consumption	
	kW	kW	% of Whole Building	kW	% of Whole Building
A	2990	1500	51	1010	34
B	3900	2040	51	1270	32
C	1760	1110	63	442	25
D	4742	1876	40	942	20

Sources: PG&E Study and LBNL Benchmark results

Load Intensity



Sources: PG&E study by Rumsey Engineers and LBNL benchmark results

Are Design Loads Overstated?

Table 3. Diversity Factor of Server Load Energy Density

	Designed Server Load (W/sf)	Extrapolated Full Occupancy Server Load (W/sf)	Server Load Diversity Factor
Facility A	50	32	0.64
Facility B	65	52	0.80
Facility C	75	52	0.69

Source: PG&E Study by Rumsey Engineers

Efficiency Opportunities

(from case studies to date)

- ◆ Raise chilled water temperature
- ◆ Turn off CRAC units in unused areas
- ◆ Use of air-side economizers
- ◆ Eliminate air-cooled chillers


More Efficiency Opportunities

(from case studies to date)


- ◆ Use evaporative cooling
- ◆ Eliminate secondary chilled water loops
- ◆ Use variable speed drives on secondary chilled water loop
- ◆ More precise air distribution

What Does The Future Hold?

- ◆ Processor and storage power (and heat that must be removed) has been increasing
- ◆ More computing power in smaller spaces
- ◆ Some technologies improve energy performance (i.e. copper)
- ◆ Some fear air cooling limits are being reached
- ◆ Does the average heat load increase or decrease?




Product Heat Density Trend Chart




Source: Uptime Institute

Theoretical Center




Source: Liebert Corp.

Impact on Raised Floor Height



Source: Liebert Corp.

Temperature Rises Quickly




Source: Liebert Corp.

Economic and Reliability Reality

- ◆ At What Point Do the Economics of Real Estate and Cooling Override the Desire to Compress Loads into Denser Spacing That Require Alternate Cooling Strategies?
- ◆ Alternative Complex Cooling Strategies Will Add Failure Modes (or ability to recover) Thereby Decreasing Reliability.

Why a Data Center Research Roadmap?

- ◆ Data Centers are Important to California Economy
- ◆ Concern over Demand and Distribution
- ◆ California wants the “Industry” to Identify Public Goods Research Needs and Priority
- ◆ Roadmap Will Aid in Coordination of Research and Deployment
- ◆ Roadmap May Stimulate Other Industry Action
 - Uptime Institute
 - ASHRAE
 - iTherm
 - CEETHERM
 - 7X24 Exchange ?



Roadmap Topics

- ◆ Market Issues
- ◆ Benchmarking (Energy intensity, end use, system efficiency)
- ◆ Tools (Energy modeling; CFD analysis; etc.)
- ◆ Building Systems
 - HVAC (Chilled water efficiency, CRAC efficiency, etc.)
 - Controls
 - Electrical (Distribution, AC-DC conversion, UPS, lighting, etc.)
- ◆ Interface with IT Equipment (Rack configuration)
- ◆ Efficient Heat Removal in Electronics Equipment

Roadmap Topics, continued

- ◆ Deployment (Public Goods and industry research; How to implement new technology)
- ◆ Organizations interested in research related to data centers
- ◆ Barriers – What is blocking progress and how can it be overcome?

Brainstorming

Research and Development Needs

Where should public goods research efforts be directed?

Suggested Research Needs

- ◆ Sort through claims of blade server mfgs.
- ◆ Understand the basis of the Uptime Institutes projections
- ◆ Determine actual loads of various types of present day facilities (benchmarking)
- ◆ Research trends with more efficient processors, routers, switches, and storage devices
- ◆ Research economics and risks of direct cooling systems
- ◆ Research power supply and distribution options

Collaboration Opportunities

◆ Related Activity

- 7X24 Exchange
- iTherm
- CEETHERM
- Uptime institute
- HP and other mfgs.
- Telecom Association – USTA
- Intel Developers forum
- AFCOM (management professional)
- EPRI – Consortium for Electric Infrastructure to Support a Digital Society (CEIDS)
- SVMG
- Utilities and other public goods efforts

◆ Integration with Distributed Power initiatives

Data Center websites

- ◆ <http://Datacenters.lbl.gov> (soon)
- ◆ <http://N4E.lbl.gov/>
- ◆ <http://enduse.lbl.gov/projects/infotech.html>
- ◆ <http://enduse.lbl.gov/Info/datacenterreport.pdf>
- ◆ www.upsite.com
- ◆ www.7X24exchange.org
- ◆ www.itherm.org
- ◆ <http://www.me.gatech.edu/me/publicat/brochures/Mettl/Bro0302.htm>

**Appendix XVI - 10-16-02 Workshop meeting notice and
presentation**

October 16, 2002

7 X 24 Exchange Bay Area Chapter Meeting

Energy Efficiency Research Ideas – Roadmap input

1. Trial sites using existing Data Center locations with similar demand and expansion to create benchmark reference. Establish not only quantity but also quality of power being consumed by loads. Trend PQ & PC and cross-reference actual measurements rather than theoretical data.
2. Power Quality systems, monitoring and data collection should be built in feature of new data center designs.
3. Power Quality should be addressed at the same level as UPS or backup systems in the design stage of a data center.
4. Reduce harmonic currents via blocking filters?
5. Apply VFD's to CRAC units?
6. Apply VFD's to the fans and chillers and optimize energy use via serial communication to the site BMS.
7. Active harmonic filters/Power factor controllers parallel connected to the load bus.
8. Is there any practical limit to Watt density for future equipment?
Review issues related to "increasing the depth of the raised floor area" in an existing (older) data center. Ex. 18" raised floor with a 30 yr. Collection of type 1 and type 2 cables (as much as 75-80% of area occupied by cabling.)
9. Cooling Strategies to reduce cooling necessities are fine. Research to reduce system harmonics and the heat they impose on the distribution system.
10. Reduction of Electrical Harmonic distortion as it relates to lowering energy consumption.
11. Better & more complete data on current & future watts/sq. ft. load data.
12. Further research on distributed cooling and power matrix.
13. Efficient room height
 - a. Should we look at upward ie. higher rack and therefore higher headroom or
 - b. Should we look at sideways ie. more floor space to distribute the load?
14. Co-gen facilities in California – what's the future?
15. Floor/Power distribution planning & Design. What can we do to help design a data center when we do not know what #'s, types of equipment that's likely to be installed over the life of the data center?
16. Research would be helpful in terms of what is the efficiency of components of the data center. Are there distribution components that can make the system more efficient?
17. What are paybacks for the owners in applying the most efficient systems?
18. Low temperature cooling methodologies –
 - a. Air temp delivered @ +40F instead of +48F to +50fF at discharge point
 - b. Using industrial refrigerant techniques
 - ◆ Liquid overfeed

◆ Ammonia

◆ ???

- c. Hi-Volume air delivery techniques
 - d. System efficiencies using different combinations
 - e. Real issues of humidification/dehumidification
19. Study effects of parasitic loads, non-linear influence & harmonic rich environments on supply-side load delivery & management. Develop life cycle costing vs. initial capital outlay curves as related to power delivery/power density concerns.
 20. Power supplies internal to servers efficiency & reductions in losses
 21. Overhead cooling vs. underfloor
 22. UPS "Battery vs. rotary
 23. Server loads actual per server loads – lists by manufacturer comparing actual vs. max loads
 24. Open racks vs closed racks w/doors
 25. Realistic max. rack loads for raised floor cooling
 26. Efficiency of architectural volume or spatial configuration & construction building materials
 27. Spot cooling strategies for loads exceeding 1500 W/SF in server cabinets – or 9600 W/rack
 28. Optimal sizing of server room for air cooling
 29. Server cooling designs front to back airflow vs. top to bottom airflow for cabinets.
 30. Use of low power chipsets and internal system cooling
 31. Develop standardized methodology for air flow within servers
 32. Somewhere there comes a trade-off in the high density data center between IT efficiency and the desire to provide cooling or heat removal with air.
 33. I'd like to see data on how much power is consumed per device. In other words, in order to plan the watts/sq.ft. consumed by a set up of various devices we need good data (not nameplate amp draw).
 34. I'd like to use a computer modeling program to predict airflow above the raised floor. I know they exist for below floor modeling.

Appendix XVII - 4-17-03 Meeting notice and presentation

California Data Center Research

Sponsored by:
California Energy Commission

Bill Tschudi
Lawrence Berkeley National
Laboratory



Acknowledgements

- ◆ California Energy Commission
- ◆ Pacific Gas and Electric Company
- ◆ 7 x 24 Exchange
- ◆ NYSERDA
- ◆ E Source
- ◆ Rumsey Engineers
- ◆ Uptime Institute
- ◆ RMI
- ◆ Industry Partners (Too many to name all)

AGENDA

9:00 – 9:15	Introductory Remarks
9:15 – 9:45	Report on LBNL research
9:45 – 10:00	Rumsey Engineers - Case Study report
10:00 – 10:30	Overview of Roadmap and draft priorities
10:30 – 10:45	Instructions for Breakout Session
10:45 – 11:00	Break
11:00 – 12:00	Breakout – Prioritize Activities
12:00 – 12:45	Working lunch – Report out
12:45 – 1:00	Conclusions & Wrap up
1:00 – 3:00	7 x 24 Exchange

We were going to meet
at our NERSC facility...



California Data Center Activities

- ❖ Case Studies
- ❖ Characterize CA Market
- ❖ Energy Research Roadmap

Case Studies/Benchmarks

◆ California - 6 Data Centers at 4 Sites

- ❑ Storage Device Mfg. – 2 data centers, Sunnyvale CA
- ❑ Bank – San Francisco, CA
- ❑ Web hosting – 2 data centers - San Jose, CA
- ❑ State facility – Sacramento, CA

◆ New York - NYSERDA

- ❑ Recovery Center – Upstate New York
- ❑ Financial Institution – New York, NY

Case Studies/Benchmarks


◆ Federal Energy Management Program

- ❑ Federal facility – 2 data centers, Fresno, CA
- ❑ LBNL operated data center – Oakland, CA

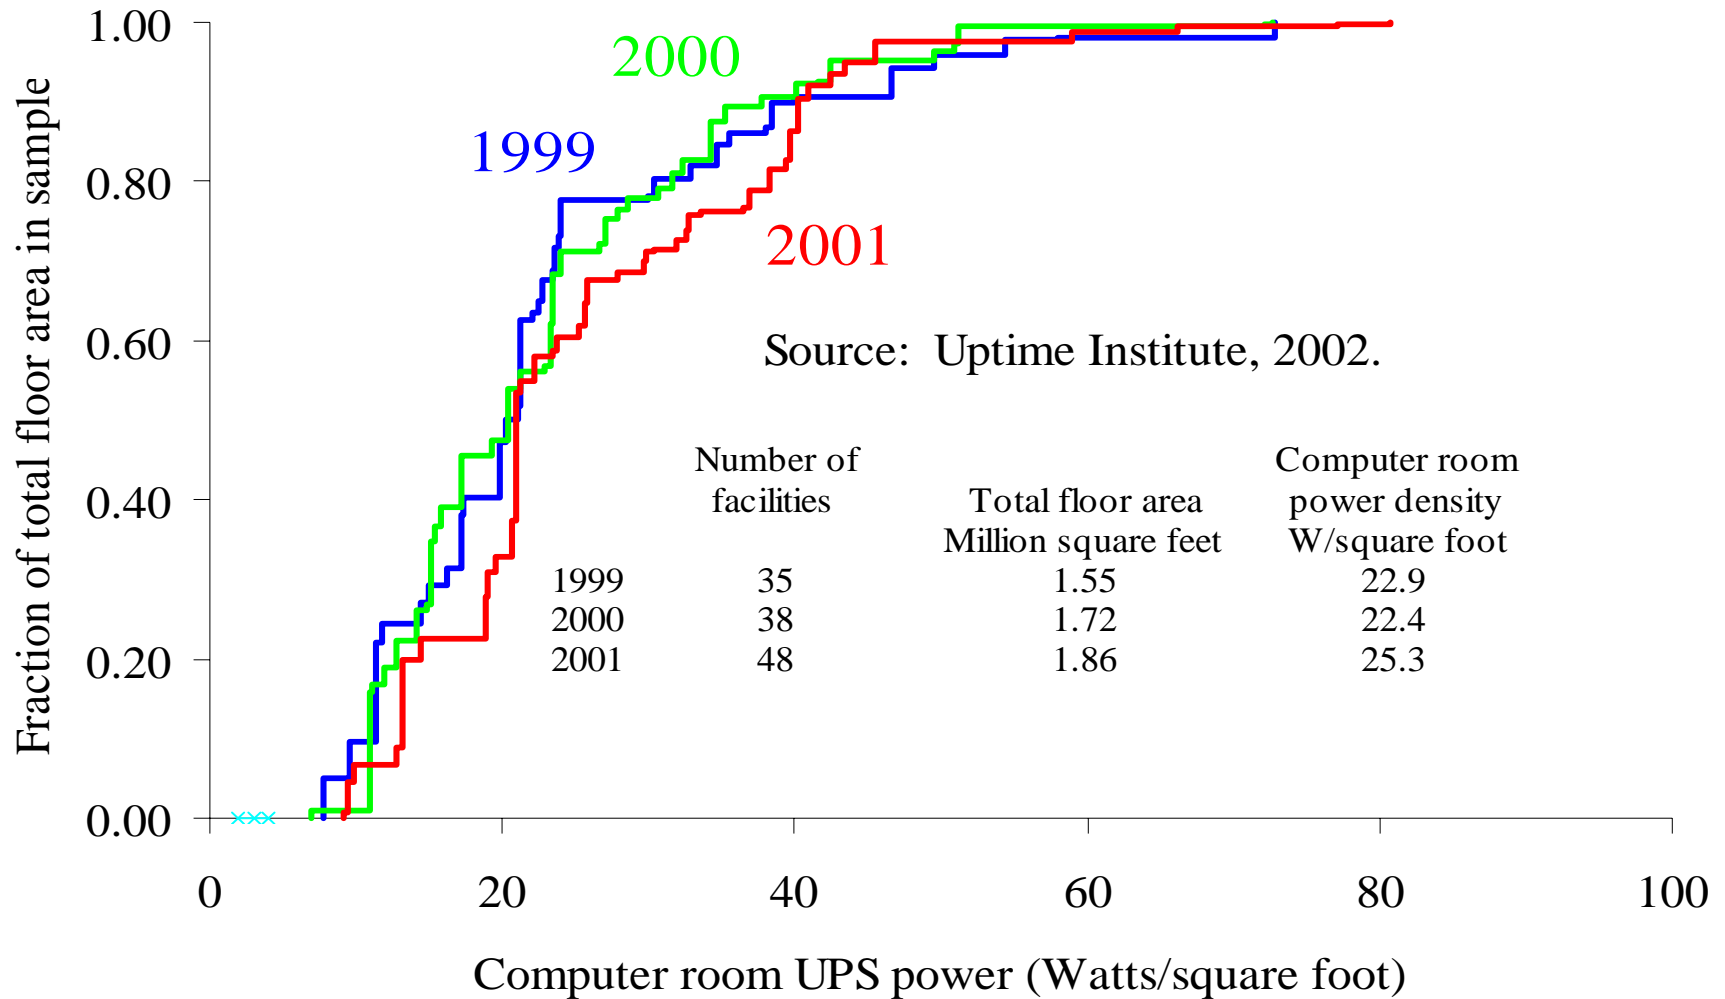
◆ Jennifer Mitchell-Jackson Case Study

Benchmarking Computer Loads

(W/Sq.Ft. of electrically active floor space)




Distribution of Computer Room Power Reported to Uptime Institute




Projecting Fully Loaded Conditions

(W/Sq.Ft. of electrically active floor space)


Current and Projected Load Intensity




Chiller Comparison




Total Chilled Water System Efficiency




Loss in UPS as a percent of total



HVAC (as a % of total load)



Better Ratio



Index of Performance

The Uptime Institute proposes a metric termed:

Index of Performance = Building systems KW ÷ UPS Output

Metrics

◆ Computer load intensity:

W/sq. ft. electrically active space

◆ UPS losses: %

◆ Chilled water: kW/ton; W/sq.ft.

◆ End use pie chart: W/end use; W/sq.ft.

◆ Occupancy:

◆ % full (subjective)

◆ % loaded – UPS or PDU

Possible Additional Benchmarks

- ◆ Computations per Watt
- ◆ Nameplate vs. Actual Comparisons
 - ◆ IT Equipment
 - ◆ UPS
 - ◆ Chillers
 - ◆ Transformers
- ◆ Standby generator energy losses
- ◆ Others?

California Data Center Market

◆ Little published information

- ❑ 2000 Solomon Smith Barney (revised downward by 25% in 2001)
- ❑ 2001 Robertson Stevens

◆ 37 centers – SF/San Jose

26 centers – LA/San Diego

4 centers - Sacramento

California Data Center Market

Grubb and Ellis figures published in the Wall Street Journal in Jan, 2003:

San Francisco/San Jose: 5.4 million sq. ft.

Los Angeles: 3.8 million sq. ft.

Includes spaces larger than 20,000 sq. ft. and at least 75% dedicated to “data center”

California Data Center Market

◆ Preliminary Floor Area Estimates

- ◆ Hosting Facilities 2 million sq.ft.
- ◆ Corporate Facilities 2-6 million sq. ft.
- ◆ Institutions 2-6 million sq.ft.
- ◆ Total Estimate 6-14 million sq.ft.

◆ Preliminary Electrical Estimates

300-700MW or 2.55 – 5.95 TWh/year

April 10, 2003 San Jose Mercury News

“A new power plant is up and running in San Jose's Alviso neighborhood, but the massive Internet server farm that it was supposed to fuel is nowhere in sight.

The Los Esteros Critical Energy Facility, a 180-megawatt plant built by Calpine in North San Jose, was designed to power an adjacent Internet server farm by U.S. Dataport. The server farm never broke ground -- and company officials didn't return calls Wednesday to say if or when it might -- but Calpine proceeded with the plant anyway, after securing a three-year deal with the state Department of Water Resources to buy power.

Company and state officials say the plant is still needed, even though the state's infamous energy crunch of 2000-01 is long over.”

180 MW: 900,000 sq.ft. x 200 W/sq.ft

California Case Studies

Our sub-consultant, Rumsey Engineers, performed data collection, preliminary analysis, provided efficiency recommendations, and case studies reports.

Data Center Energy Roadmap

- ❑ Input through workshops, conferences, and contacts with Industry professionals.
- ❑ Participation in design charrette facilitated by the Rocky Mountain Research Institute (RMI)
- ❑ Additional input today
- ❑ Goal for today – CONFIRM PRIORITIZATION

Summary

Activities	Priority 1 Highest	Priority 2	Priority 3	P/I
UNDERSTAND THE DATA CENTER MARKET				
Market Characterization				
Monitor utility load requests			X	P
Refine and update data center market assessment <u>2M or 20M square feet</u>		X		P
Monitor trends in processor heat intensity	<u>X</u>		X	<u>I</u> /P
<u>Greening data centers</u> <u>????</u>			<u>X</u>	
ENERGY BENCHMARKING				
Perform additional case studies and benchmarking	X			P
Encourage sub-metering to facilitate end use benchmarking		X		P
Compare IT equipment measured vs. nameplate	X			P/I
Develop self-benchmarking protocol	X		<u>X</u>	P
Develop and maintain benchmarking database	X	<u>X</u>		P
Develop additional metrics <u>(include below)</u>		X	<u>X</u>	P/I
<u>Standardize metrics</u> <u>???</u>	<u>X</u>			
<u>Benchmark new technologies</u>			<u>X</u>	
IDENTIFY AND DEVELOP BEST PRACTICES				
Determine Current Best Practices				
Document best practices observed in case studies	X			P
Research available modeling tools and provide designers with comparative data		X		P
Document current operating temperature and humidity set points to determine widest acceptable ranges		X		P/I
Survey available energy storage technologies	X			P
<u>Link reliability with energy efficiency</u>				
Improving Best Practices				
Link improved availability, reliability, and energy efficiency	X			P
Develop efficient part load and incremental build-out strategies	X			P
Develop facility planning and design guides	X			P
Develop and demonstrate air management improvements	X			P/I
Develop UPS sizing guidelines including life cycle cost evaluation	X			P
Encourage increased use of air and water economizers – free		X		P/I

cooling				
Encourage more efficient drives including variable speed on CRAC units and on conventional air handlers		X		I
<u>Align incentives</u>				
IMPROVING INFRASTRUCTURE SYSTEMS EFFICIENCY				
General				
Develop advanced design and modeling tools		X		P
Manage computing load to avoid local "hot spots" -controlled by improved temperature monitoring (<u>concern with encouraging more computing power than needed</u>)	X			P/I
Develop and implement modular (scalable) system concepts to improve part load efficiency		X		P/I
Develop mechanical and electrical system sizing guidelines including use of benchmark results to account for load diversity	X			P
HVAC				
Improve energy efficiency of CRAC units (e.g. variable speed compressors, improved controls, etc.)		X		I
Develop improvements in data center air distribution systems (<u>may be included below</u>)	X			P/I
Increase allowable ranges of temperature and humidity		X		I
<u>Optimize air movement/heat removal (through room, racks, equipment, and back)</u>	<u>X</u>			
<u>Improve part load efficiency to overcome impacts Over sizing</u>	<u>X</u>			
<u>Improve Controls and control integration</u>	<u>X</u>			
<u>Increase water and air side economizers</u>		<u>X</u>		
Electrical				
Improve UPS efficiency, especially at part load, through testing and publication of UPS efficiency vs. load	X			P
Develop more efficient solutions for given levels of reliability (N+1, N+2, 2N, etc.) with UPS systems		X		P/I
Improve energy storage technologies for data centers		X		I
Assess opportunities to improve efficiency through optimised distribution such as minimizing conversions and DC		X		P

distribution <u>(identify missing pieces)</u>					
Encourage efficient lighting – efficient systems, lighting controls, use of maintenance task lighting			X		P/I
<u>Implementation protocol,(design through to operations)</u>	<u>X</u>				<u>P/I</u>

INTERFACE OF BUILDING SYSTEMS AND IT EQUIPMENT				
Encourage interaction of IT and facilities professionals to specify energy efficient racks and cabinets		X		I
Investigate limits of air cooling leading to direct cooling (liquid) rack and computer systems	<u>X</u>	X		I
Evaluate trade-off of increasing density vs. spreading out heat load. Develop economic evaluation model.	<u>X</u>	X		P
Evaluate emerging rack and cabinet systems for energy efficiency	X	<u>X</u>		P
<u>Decrease equipment heat</u>	<u>X</u>			<u>I</u>
COMPUTER (IT EQUIPMENT)				
Investigate more efficient power supplies through assessment of current practice and available technology	X	<u>X</u>		<u>I/P</u>
Work with industry to determine strategies to improve power supply efficiency (e.g. labelling and rebate programs)	X			P/I
Monitor trends in heat intensity and cooling options		X		<u>I/P</u>
Investigate high efficiency cooling such as spray cooling to remove heat from processors <u>(integrate with above)</u>		X		<u>P/I</u>
Benchmark emerging technologies (such as Blade Servers) for Energy Performance			X	I
Manage computing load and demand through improved <u>power management</u> <u>"sleep modes"</u> <u>(not clear)</u>			X	I
<u>Document relative efficiency of chips</u>				
<u>Box drives the system – define energy "U" develop efficiency metrics for IT equipment</u>	<u>X</u>			

Consider greater emphasis on controls

Increase emphasis on systems approach (integration)

Develop standards as well as guides

Establish champion industry group (e.g. energy efficiency subgroup of 7x24)

Don't forget back of house, secondary areas

Alternative technologies to cool electronics

Reduce Operating costs

Energy Efficiency and Reliability

Improve ROI

Improved reliability

No impact to IT organizations

Improvement

No/low cost impact

Does it cross boundaries (barriers)

Design for expandability – modular

Normalized metrics and processes

Public information / Justification for public funding

Breakout Schedule:

Start - 11:20

Review Activity list – 10 minutes 11:30

Review and Prioritize activities – 10 minutes 11:40

Add activities – As Great ideas surface

Fill out forms/designate spokesperson – 10 minutes 11:50

Float – 10 minutes 12:00

Get Lunch (Thanks to PG&E and 7x24) – 10 minutes 12:10

Report outs – 5 minutes each group (30 minutes) 12:40

Wrap – 10 minutes 12:50

Float – 10 minutes 1:00

Summary

Activities	Priority 1 Highest	Priority 2	Priority 3	P/I
UNDERSTAND THE DATA CENTER MARKET				
Market Characterization				
Monitor utility load requests			X	P
Refine and update data center market assessment <u>2M or 20M square feet</u>		X		P
Monitor trends in processor heat intensity	<u>X</u>		X	<u>I</u> /P
<u>Greening data centers</u> ????			<u>X</u>	
ENERGY BENCHMARKING				
Perform additional case studies and benchmarking	X			P
Encourage sub-metering to facilitate end use benchmarking		X		P
Compare IT equipment measured vs. nameplate	X			P/I
Develop self-benchmarking protocol	X		<u>X</u>	P
Develop and maintain benchmarking database	X	<u>X</u>		P
Develop additional metrics <u>(include below)</u>		X	<u>X</u>	P/I
<u>Standardize metrics</u> ???	<u>X</u>			
<u>Benchmark new technologies</u>			<u>X</u>	
IDENTIFY AND DEVELOP BEST PRACTICES				
Determine Current Best Practices				
Document best practices observed in case studies	X			P
Research available modeling tools and provide designers with comparative data		X		P
Document current operating temperature and humidity set points to determine widest acceptable ranges		X		P/I
Survey available energy storage technologies	X			P
<u>Link reliability with energy efficiency</u>				
Improving Best Practices				
Link improved availability, reliability, and energy efficiency	X			P
Develop efficient part load and incremental build-out strategies	X			P
Develop facility planning and design guides	X			P
Develop and demonstrate air management improvements	X			P/I
Develop UPS sizing guidelines including life cycle cost evaluation	X			P
Encourage increased use of air and water economizers – free		X		P/I

cooling				
Encourage more efficient drives including variable speed on CRAC units and on conventional air handlers		X		I
<u>Align incentives</u>				
IMPROVING INFRASTRUCTURE SYSTEMS EFFICIENCY				
General				
Develop advanced design and modeling tools		X		P
Manage computing load to avoid local "hot spots" -controlled by improved temperature monitoring (<u>concern with encouraging more computing power than needed</u>)	X			P/I
Develop and implement modular (scalable) system concepts to improve part load efficiency		X		P/I
Develop mechanical and electrical system sizing guidelines including use of benchmark results to account for load diversity	X			P
HVAC				
Improve energy efficiency of CRAC units (e.g. variable speed compressors, improved controls, etc.)		X		I
Develop improvements in data center air distribution systems (<u>may be included below</u>)	X			P/I
Increase allowable ranges of temperature and humidity		X		I
<u>Optimize air movement/heat removal (through room, racks, equipment, and back)</u>	<u>X</u>			
<u>Improve part load efficiency to overcome impacts Over sizing</u>	<u>X</u>			
<u>Improve Controls and control integration</u>	<u>X</u>			
<u>Increase water and air side economizers</u>		<u>X</u>		
Electrical				
Improve UPS efficiency, especially at part load, through testing and publication of UPS efficiency vs. load	X			P
Develop more efficient solutions for given levels of reliability (N+1, N+2, 2N, etc.) with UPS systems		X		P/I
Improve energy storage technologies for data centers		X		I
Assess opportunities to improve efficiency through optimised distribution such as minimizing conversions and DC		X		P

distribution <u>(identify missing pieces)</u>					
Encourage efficient lighting – efficient systems, lighting controls, use of maintenance task lighting			X		P/I
<u>Implementation protocol,(design through to operations)</u>	<u>X</u>				<u>P/I</u>

INTERFACE OF BUILDING SYSTEMS AND IT EQUIPMENT				
Encourage interaction of IT and facilities professionals to specify energy efficient racks and cabinets		X		I
Investigate limits of air cooling leading to direct cooling (liquid) rack and computer systems	<u>X</u>	X		I
Evaluate trade-off of increasing density vs. spreading out heat load. Develop economic evaluation model.	<u>X</u>	X		P
Evaluate emerging rack and cabinet systems for energy efficiency	X	<u>X</u>		P
<u>Decrease equipment heat</u>	<u>X</u>			<u>I</u>
COMPUTER (IT EQUIPMENT)				
Investigate more efficient power supplies through assessment of current practice and available technology	X	<u>X</u>		<u>I/P</u>
Work with industry to determine strategies to improve power supply efficiency (e.g. labelling and rebate programs)	X			P/I
Monitor trends in heat intensity and cooling options		X		<u>I/P</u>
Investigate high efficiency cooling such as spray cooling to remove heat from processors <u>(integrate with above)</u>		X		<u>P/I</u>
Benchmark emerging technologies (such as Blade Servers) for Energy Performance			X	I
Manage computing load and demand through improved <u>power management</u> <u>"sleep modes"</u> <u>(not clear)</u>			X	I
<u>Document relative efficiency of chips</u>				
<u>Box drives the system – define energy "U" develop efficiency metrics for IT equipment</u>	<u>X</u>			

Consider greater emphasis on controls

Increase emphasis on systems approach (integration)

Develop standards as well as guides

Establish champion industry group (e.g. energy efficiency subgroup of 7x24)

Don't forget back of house, secondary areas

Alternative technologies to cool electronics

Reduce Operating costs

Energy Efficiency and Reliability

Improve ROI

Improved reliability

No impact to IT organizations

Improvement

No/low cost impact

Does it cross boundaries (barriers)

Design for expandability – modular

Normalized metrics and processes

Public information / Justification for public funding

Breakout Schedule:

Start - 11:20

Review Activity list – 10 minutes 11:30

Review and Prioritize activities – 10 minutes 11:40

Add activities – As Great ideas surface

Fill out forms/designate spokesperson – 10 minutes 11:50

Float – 10 minutes 12:00

Get Lunch (Thanks to PG&E and 7x24) – 10 minutes 12:10

Report outs – 5 minutes each group (30 minutes) 12:40

Wrap – 10 minutes 12:50

Float – 10 minutes 1:00



Bay Area Chapter

www.7x24sf.org

www.7x24exchange.org

Spring Meeting Details:

When: Thursday, April 17, 2003
9:00am – 2:00pm
** Working Lunch provided by PG&E and 7x24 Exchange **

Where: PG&E Energy Center
851 Howard Street, San Francisco
(Directions and contact information can be found at the following link:
http://www.pge.com/003_save_energy/003c_edu_train/pec/admin/contact.shtml)

RSVP: 7x24rsvp@mazzetti.com by April 14th, 5:00pm.
Please respond by deadline as lunch is being provided.

Agenda:

◆ Data Center Power, Bill Tschudi, Lawrence Berkeley National Labs

LBNL has been the chief advocate for data center energy consumption research in large- and small-scale data centers. This has a profound affect on the size and type of facilities we design and operate. Bill will present a draft findings of his report, and he's seeking YOUR input as to it's content. This is a unique opportunity for you to view and comment on technology's impact on our facilities before this report's released. We all expect an energetic discussion on platform and facility W/sf!


A link to Lab's previous work is: <http://datacenters.lbl.gov/>. The draft report will be presented in this meeting for your commentary and discussion.

◆ Data Center Planning, Doug Schmitt, Jones Lang LaSalle

Doug has recently returned from giving a 5-hour tutorial on Data Center Best Practices at the recent AFCOM conference. The presentation was well received, and Doug will be covering several topics on data center planning and practices such as security, process management and project facilitation. Several "real world" situations will be presented and discussed.

Facility 8 in Perspective

HVAC Power Consumption



Facility 8 Site Characteristics

Data Center 8.1

- 26,200 sq ft
- 6 UPS's – 3 per “side”
- Redundancy: n+1 at PDU level, n+2 at UPS level
- Overhead ducted air distribution
- Air-cooled constant volume CRAC units




Facility 8 Site Characteristics

Data Center 8.2

- 73,000 sq ft
- 5 UPS's
- Redundancy: n+1 at PDU level
- Overhead ducted air distribution
- Central Chilled Water Plant
- Central air handling system
- Variable speed chiller,
secondary pumps, air handlers




Whole Building End Use




Facility 8 Electricity End-Use

Data Center 8.1



Data Center 8.2



Efficiency Metrics

Data Center 8.1

Metric	Value	Units
Data Center Computer Power Density	8.5	W/sf
Data Center Cooling Power Density	12	W/sf
Cooling kW : Computer Load kW	1.4	--

Data Center 8.2


Metric	Value	Units
Data Center Computer Power Load Density	14.5	W/sf
Data Center Cooling Power Load Density	5.3	W/sf
Cooling kW : Computer Load kW	0.4	--
Chiller Efficiency	0.4	kW/ton
Chilled Water Plant Efficiency	0.6	kW/ton

Cooling Efficiency - Other Data Centers

System Type	Efficiency (kW/kW)
Air cooled chillers and fan coil units	0.5
Air cooled CRAC units	0.8
Water cooled reciprocating chiller and CRAC units	1.5
Water cooled centrifugal chiller plant with CRAC units	0.7


Chiller Efficiency

Chiller Efficiency



Chilled Water Plant Efficiency

Chilled Water Plant Efficiency



How did they do it?

Data Center 8.1

- Air cooled CRAC units
- No economizing
- Constant Speed Fans
- Humidification control
- All CRAC units on

Data Center 8.2

- Optimal central chilled water plant
- Optimal central air handling units
- Little humidity control
- Good control
- Data Monitoring – Gateways, EMCS

How can they do even better?

Data Center 8.1 Recommendations

- Disable humidification control
- Turn off CRAC units
- Rotate UPSs
- Space temperature setpoints



How can they do even better?

Data Center 8.2 Recommendations

- Monitoring - chiller, total chiller plant kW/Ton
- Run Cooling towers in parallel, nozzle replacement
- Chilled water setpoint
- Condenser water temperature reset



Common Findings

- ❑ Humidity Control and CRAC unit fighting
- ❑ Lighting Control
- ❑ Air side Economizing
- ❑ Variable Speed Drives – pumps, chillers, fans
- ❑ Control strategies - setpoints, cooling tower staging
- ❑ High ceilings important with Underfloor (thermal stratification)
- ❑ Air management common problem
 - ❑ rack configuration, underfloor vs. overhead, underfloor congestion
- ❑ Low UPS loading

**Appendix XVIII – ACEEE 2003 paper #162, “Data Centers and
Energy Use – Let’s Look at the Data”**

Data Centers and Energy Use – Let's Look at the Data

William Tschudi, Lawrence Berkeley National Laboratory

Priya Sreedharan, Rumsey Engineers, Inc.

Tengfang Xu, Lawrence Berkeley National Laboratory

David Coup, NYSERDA

Paul Roggensack, California Energy Commission

ABSTRACT

Data centers are prevalent in a wide range of industries, universities, and government facilities. Energy demand in these facilities is thought to be growing as computing technology changes and IT professionals seek to maximize computing per square foot of data center. In addition, a multitude of methods have been used to estimate and quantify energy intensity. As a result, there is considerable confusion over data centers' electrical use today, and the needs for the future. Research aimed at understanding the present electrical intensity, end use, and key facility systems' operation was undertaken with the ultimate goal of developing energy efficiency improvements in these buildings.

Metrics used in this study allow comparison of the current power density of computing equipment, and provide an indicator of the efficiency of key facility systems. In addition, a metric is included to evaluate how efficiently the Heating, Ventilating and Air Conditioning (HVAC) system operates to maintain satisfactory conditions for computing equipment. A review of the summary benchmark results can identify more efficient systems and practices, and can discover operational problems.

The information presented can be used by data center owners, operators, and designers to understand current performance, to set design and operational criteria for new projects, to identify current best practices in design and operation, and to improve reliability. This study also identifies gaps where additional research is needed to achieve a new level of improved energy efficiency.

Introduction

Data Centers are prevalent and an integral part of today's world. Data centers support the communications industries, and play a major role in the Internet economy. They also support research and learning, and are the "brains" within most corporations and government institutions. The California Energy Commission and the New York State Research and Development Authority realize the importance of these facilities and the opportunity to improve their energy efficiency. Although electrical power demand is high, designers and operators of data centers currently have little information concerning where to place their resources to improve their efficiency. In addition, there is little benchmark data available to highlight what can be achieved in the design of new systems.

The benchmarking reported here involved a strategy that obtained an energy end use breakdown in a number of data centers. A broad definition of a data center was adopted, since similar building issues are present regardless of the computing platform. Two key objectives led to this work: First, utilities, public interest organizations, and those that work with data centers, all have a critical need for more information concerning current data center electric power requirements, as well as future energy demand trends. Energy benchmarking sets the stage for improvements by documenting current energy use and intensity, and over time, can be trended to establish further guidance. The second objective was to identify the energy end-use in data centers. Then, efficiency opportunities could be targeted for each of the intensive areas.

Definitions

Acronym	Definition
HVAC	Heating, Ventilation, and Air Conditioning
UPS	Uninterruptible Power Supply
CRAC unit	Computer Room Air Conditioning unit
IT	Information Technology
VAV	Variable Air Volume
PDU	Power Distribution Unit
VFD	Variable Frequency Drive

Background

Energy demand of today's IT equipment and electrical power for systems removing the heat they produce are high compared to ordinary commercial buildings. Although data centers contain various types of computing equipment, building systems in data centers usually have similar characteristics and can account for more than 50% of the total energy. A segment of the data center market even uses excess infrastructure as a selling point, resulting in oversized electrical and HVAC systems. Unfortunately, these oversized systems usually operate inefficiently.

Planning for the future, whether at the utility level, facility level, or computing equipment level has been a challenge. Previously, little publicly available energy benchmark data existed for data centers. Since confidentiality of facility operating information is important to a majority of data center operators, reliable building energy benchmark information was not made available. Data center operators typically track whole building energy use and energy used by the computing equipment. What has been lacking is measured data for comparison of electric power density (Watts/sq ft), energy end use, and efficiency comparisons of key facility systems (HVAC, UPS, lighting, etc.).

HVAC systems typically include computer room air conditioning and ventilation, and may include a large central cooling plant. In addition, lighting and other minor loads are present. Reported here are the energy benchmark results, but since data centers typically are large heat sources and generate the need for high

cooling loads with tightly controlled environmental conditions, much of the study focuses on the efficiency of the HVAC systems.

Buildings with data centers have large, constant electrical demand to operate the computing equipment. Current technology has evolved to a practice where computing equipment is typically air cooled through use of energy intensive HVAC systems consisting of large central plant heating and cooling, and use of computer room air-conditioning (CRAC) units or other large air handlers. Data centers often mandate strict environmental design considerations calling for tightly controlled temperature and humidity with the objective of protecting the computing equipment from overheating.

Benchmarking Activity

A measurement methodology and metrics most useful for comparing data centers and their facility systems was developed. The metrics allow comparison of widely varying data centers regardless of the design, and the types of computing equipment. These metrics illustrate measured electric intensity, which is useful to trend overall load growth and to predict future needs. They also provide insight into how efficiently the building systems were designed and are operating. Energy use and systems operational information was obtained primarily on chillers, UPS systems, and CRAC units. This data was obtained by connecting power sensors to the host electrical panels, or by reading from the equipment's digital meters, if in existence. Additional operational data, such as flow, and temperature measurements were obtained from existing facility management systems to the extent they were available and finally were supplemented by direct measurement if not readily available.

Tables 1 and 2 summarize the metrics and other information used in this study. Ten data centers in various industries, housing various types of computing equipment were included in this study.

Table 1. Data Center Metrics

Whole Building Electrical Power:	kW
Load Intensity:	
Data Center floor area	square feet (sq ft)
Total load density	W/ sq ft
Computing load density	W/ sq ft
HVAC load density	W/ sq ft
HVAC:	
Chiller plant	
Chiller Efficiency	kW/Ton
Chilled Water Plant Efficiency	kW/Ton
Chiller load	Tons
Data Center Load	Tons
HVAC air systems	
CRAC unit fan, and humidity control energy	
Central air handling fan power	Cubic Feet per Minute
Air handler fan efficiency (where possible to obtain)	per kW (CFM/kW)

External temperature and humidity	°F, %
Data Center Electrical power demand:	
UPS Loss	kW
Computer load (from UPS Power)	kW
HVAC - chilled water plant (if central plant exists)	kW
HVAC - central air handling, and/ or CRAC Unit energy	kW
Lighting	kW
Design Data:	
Design basis for Computer load	kW/sq ft
Design basis for Chilled Water, air side HVAC, and UPS Systems	Temperature
	Humidity
	Flowrate
	% Efficiency
	Total load
	etc.

Table 2. Additional Data Center Information

Features and System Descriptions	Example Descriptions
HVAC	Central water-cooled chilled water plant, central air handling system with VAV control
	Distributed air-cooled CRAC units
	Air-cooled chillers with CRAC units supplying air under floor
	Central air handlers use outside air economizers
Variable-speed-drives	Centrifugal chiller with VFD
	Primary/ Secondary with VFD
	Central air handler with VFD
Electrical Distribution	N+1 UPS's
	N+1 at the PDU level
	Backup power generators
Control Strategies	Multiple cooling towers operated in parallel
	Minimum number of chillers operated
	CRAC units in empty areas turned off.
	Humidity control disabled on CRAC units
	VAV system with duct static pressure of 0.75"
	Chilled water setpoint fixed at 50 °F
	Condenser water setpoint fixed at 70 °F
	Chiller kW/Ton monitored continually
Temperature and Humidity Setpoints	Air side economizers used on
	Return air temperature maintained at 70 °F ± 5 °F
	Supply air temperature of central air handlers maintained at 50 °F
	Relative humidity maintained at 50 % ± 10%

Features and System Descriptions	Example Descriptions
Redundancy/Reliability	N+1 at UPS level
	N+1 at PDU level
Estimate of Occupancy	Data center is 40% full - physical capacity
	Operating at 30% of UPS capacity

The case studies and summary benchmarking data are available through the LBNL website: www.datacenters.lbl.gov


During the project, the on-site team noted potential efficiency opportunities through visual observation, analysis of the data, and discussion with facility personnel. These opportunities were described to the participating facility in a final report. The observations were qualitative in nature, and were based upon the site team's prior experience and limited observations. In some cases, recommendations for further investigation were made. These recommendations typically required additional evaluation by the owner but could result in short or long-term efficiency improvement.

Fourteen data centers in eleven facilities, (where three facilities had two data centers each) were included in this study. To develop a more robust data set, many more benchmarks will be needed. Once this information is available however, building operators will be able to gauge the relative performance of their facility systems and intensities of various computing equipment. In the future, a mechanism for self-evaluation is proposed for development that would allow a data center owner/operator to compare his data center's performance to a larger sampling of data centers. This information should improve the ability to predict future power requirements and size systems more efficiently.

Benchmark Results


End use energy breakdowns were obtained for the data centers in this study. Figure 1 illustrates the energy end-use information that is provided by the benchmark measurements for a typical data center.

Figure 1. Representative Energy End Use Breakdown



Typically, the end use components consisted of the electrical loads for the computing equipment (fed from UPS systems), UPS system losses, HVAC – chilled water plant, HVAC – computer room air conditioners, and lighting. The relative percentages of each of these components varied according to the computing load intensity and the efficiency of the infrastructure systems necessary to support the computing. For example, the percentage of the total power to the computing equipment varied between 33% and 73 %. Similarly, the other end use components varied considerably as shown below in figure 2.

Figure2. Benchmarking Examples



The data center shown to the left utilized a highly efficient system that was thoughtfully designed using best practices with better than standard HVAC

components and controls. The center represented by the chart on the right utilized traditional distributed air-cooled computer room air conditioners.

In similar fashion, issues with UPS systems, lighting, and other systems are highlighted. In one facility, benchmarking discovered that the entire cooling for the computing equipment was being handled through the make up air (house) system, yet all of the computer room air conditioners were operating utilizing unnecessary fan energy and adding to the cooling load.


HVAC Systems

By focusing on the various HVAC systems and their components, the benchmark data reveals that energy use can vary by factors of 3 or more for systems that serve essentially the same purpose. The study utilizes an interesting metric to compare the relative efficiency of the overall HVAC system. By comparing the energy used for cooling the data center (i.e., the HVAC power in kW) to the UPS output, which should closely resemble the computer loads (in kW), an indicator of HVAC system performance is obtained. A lower value indicates that the system is likely to be more energy efficient. This metric is defined as follows:

$$HVACperformanceindex(\%) = \frac{kW_{HVAC}}{kW_{UPSOutput}}$$

For this study, many different HVAC system designs were observed and measured. Figure 3 shows a comparison of data for 13 system configurations in the data centers measured in this study.¹ This information highlights that there is wide variation in system design and energy efficiency. This wide variation underscores the need to understand the features and principles of the more efficient systems. This will lead to best practices in design and construction of these systems.

Figure 3. Relative Data Center HVAC System Performance



The HVAC system was further benchmarked by examining chilled water and computer room air conditioning systems. For the chilled water plant, the traditional

¹ Due to measurement complications, comparable data was not available for Facility 13.


efficiency metric, kW/Ton, was used. Chiller, pumping, and fan energy (for cooling towers) and the corresponding tons of chilled water produced were obtained. Wide variations in efficiency (dominated by the chiller efficiency) were observed.

Computer room air conditioning energy was similarly determined. Here, it is more difficult to obtain accurate airflow measurements – typically delivered by many air handlers into common underfloor areas or a network of ducting. As a result, an efficiency metric such as cfm/kW was not generally determined. We instead rely on the comparison of overall HVAC performance. This is an area where further research could pinpoint additional efficiency issues.

Computing Loads

The electrical load that the computing equipment requires must all be removed as heat by the HVAC system. The electric power density can vary significantly from data center to data center. Figure 4 shows the measured electric power density due to the computing equipment alone. In the calculation to determine this metric, the Uptime Institute’s definition of “electrically active” floor area is used in the denominator. This effectively excludes areas such as walkways or storage spaces, which are more likely to have electric power density similar to commercial office buildings.

Figure 4. Benchmarked Computer Load Densities



Uncertainties in predicting the computing equipment loads make HVAC system sizing a challenge. Measured electric power density is typically quite less than what is specified during design. This occurs for several reasons: facility designers using name plate data (which typically leads to loads that are several times greater than actual), uncertainty for future equipment power density, unnecessary or compounding conservatisms, computing equipment (and load) is added gradually over time as business needs dictate. In addition, reliability strategies require that multiple, redundant equipment be available. This introduces significant inefficiency for standby or part load operation. While the computing load varies for each

application, measured data from facilities with similar computing missions will help “right-size” the cooling equipment. Cooling systems are often more efficient when operated near their full design load. This study found computing loads at all facilities below 65 W/sq ft. Use of benchmark data can lead to better prediction of design loads and better build-out strategies. Designing systems and components in closer alignment with actual operating loads will also lead to more efficient operation.

Observed Efficiency Improvement Opportunities

Based upon the limited data collected and site observations, a number of efficiency recommendations are emerging as better practices for the facilities monitored. These are categorized in broad categories, and include design and control issues:

Table 3. Observed Energy Efficiency Opportunities

Chilled Water Plant Design Issues	<ul style="list-style-type: none"> • Chiller efficiency dominates the efficiency of the chilled water plant but this includes use of variable speed drives, raising chilled water temperature, and lowering condenser water temperature • Water cooled chillers are more efficient • Use of free cooling using cooling tower water when conditions permit. • Improve cooling tower efficiency by operating all cooling towers at reduced fan speed (variable) rather than operating fewer towers at full speed. • Reduce pumping energy through use of variable speed drives.
Air System Design Issues	<ul style="list-style-type: none"> • Improve airflow management utilizing hot and cold aisles, closing unnecessary openings in raised floors, partitioning to direct hot and cold air, and utilizing modeling programs to optimize airflow. • Use of variable speed drives on fans • Establish broader ranges of temperature and humidity control. • Use of large air handlers with ability for air-side economizing in lieu of traditional computer room air conditioners. Include high efficiency fans, motors. • Underfloor air distribution may not be necessary. Large ducted overhead systems with directed air into hot aisles can be more efficient. Consideration of thermal stratification for supply and return is essential. • All computer room air conditioners may not need to operate to maintain conditions. Turning off some units may be possible.
Electrical System Design Issues	<ul style="list-style-type: none"> • Uninterruptible power supply efficiency should be considered. Efficiencies decrease at part load conditions. Redundancy strategies and part load operation should be considered when sizing, selecting, and operating UPS systems. • Use of conventional lighting controls such as occupancy sensors, reduced lighting levels, and/or lights-out operation should be pursued.
HVAC control strategies	<ul style="list-style-type: none"> • Chilled Water supply, condenser reset strategies. CHWS can be 50 °F. Condenser water can be based off of wet bulb temperature, or if gateways exist, differential refrigerant pressure. • In overhead systems duct static pressure can be lower than 1.5 ” w.g. One facility used 0.75” w.g. resulting in efficient fan operation. • Monitoring kW/ton of chiller and chilled water plant to see if strategies

	<p>result in savings.</p> <ul style="list-style-type: none"> • Staging of chillers: VFD centrifugals should be run in parallel, constant speed chillers should be loaded as fully as possible. • Some facilities didn't have properly working economizers. Hasn't been able to get controls vendor to fix it for years. • Use of conventional lighting controls such as occupancy sensors, reduced lighting levels, and/or lights-out operation should be pursued. • Turning off humidity control in CRAC units in many climates may be possible. Units tend to fight each other with one in humidification and one in dehumidification. Steam humidification is better, not the electric humidification in CRAC units.
Rack Configuration Issues	<ul style="list-style-type: none"> • Hot and cold aisles need to be maintained. Many facilities did not follow this protocol. Better coordination between IT and facility managers is necessary. • Rack loading is often inconsistent with some racks partially loaded, some fully loaded, and others empty. Even distribution would aid thermal performance.

Conclusion and Recommendations

Energy benchmarking results can help to visualize energy end uses in complex data center facilities. For a data center owner/operator there are a number of high value benefits. Measured energy use determined by a benchmarking program can provide a baseline for tracking energy performance over time. It can be used to better predict future needs leading to more efficient sizing of supporting facility systems. Benchmarking can also be used to prioritize where resources need to be applied to achieve improvements in energy efficiency.

Use of the metrics developed for this project provides a mechanism for comparison of facility systems and components to other data centers. This is possible even though the system design and configuration may be completely different. By analyzing the variations in the data, current better practices can be identified. The strategies and configurations resulting in the most efficient operation can then be applied to new designs or retrofit into existing facilities. Large apparent variations in the energy use of systems or components may signify design, installation, operational, or maintenance problems. Finding the reason for the discrepancy could solve on going operational or maintenance problems or correct inefficiencies originally built into the facility. For data center designers, access to actual comparison data will highlight better practices and lead to new creative energy efficient designs, and operating strategies. Future activity should be directed at developing such a database, which should ideally include both measured data, and design data. In addition, a benchmarking tool is needed such that building operators can perform their own evaluations.

This benchmarking activity illustrated that certain HVAC design strategies can be far more efficient than conventional data center cooling strategies. Designs using central air handling, free cooling, water-cooled chillers, and variable speed driven mechanical equipment were found to be more energy efficient in this study. The benchmarking also identified opportunities for optimizing existing mechanical

and electrical equipment through improved control strategies, such as staging of UPSs, increasing chilled water temperatures, and better air distribution through rearrangement of racks into hot and cold aisle configurations.

Public interest funding should be provided to assist data centers in this benchmarking process, and to provide assistance in implementing solutions that will improve energy efficiency and/or reduce peak load. Furthermore, public interest funding should be directed at developing new technologies (e.g., power supplies that are energy efficient at partial loads) or innovative strategies (e.g. quantitatively compare the energy usage associated with various strategies for managing cold air inflow and hot air discharge) that would not otherwise be developed by industry. Additional insight on important R&D topics may be found in LBNL's research roadmap for High Performance Data Centers.

References

LBNL websites:

<http://datacenters.lbl.gov>

<http://enduse.lbl.gov/Projects/InfoTech.html>

NYSERDA website:

www.nyserda.org

California Energy Commission:

http://www.energy.ca.gov/process/industry/industry_index.html

Appendix XIX – Annotated Bibliography

References

ACEEE, and CECS. 2001. Funding prospectus for "Analysis of Data Centers and their implications for energy demand". Washington, DC, American Council for an Energy Efficient Economy (ACEEE); Center for Energy and Climate Solutions (CECS). July 2001.

The paper includes an overview of data centers; discusses energy use, energy choices, and energy efficiency in data centers; potential impacts of data centers; present and future regulatory issues; and business opportunities in energy services.

Aebischer, B., R. Frischknecht, C. Genoud, A. Huser, and F. Varone. 2002a. Energy- and Eco-Efficiency of Data Centres. A study commissioned by Département de l'intérieur, de l'agriculture et de l'environnement (DIAE) and Service cantonal de l'énergie (ScanE) of the Canton of Geneva, Geneva, November 15.

The study investigates strategies and technical approaches to fostering more energy-efficient and environmentally sound planning, building and operating of data centres. It also formulate recommendations on how to integrate the findings in the legal and regulatory framework in order to handle construction permits for large energy consumers and promote energy efficiency in the economic sectors. Seventeen recommendations grouped in four topics are derived from study conclusions: Transfer of the accord into an institutionalized legal and regulatory framework; Energy-efficiency policies for all large energy consumers; Preconditions, and prerequisites; Operational design of voluntary energy policies.

Aebischer, B., R. Frischknecht, C. Genoud, and F. Varone. 2002b. Energy Efficiency Indicator for High Electric-Load Buildings. The Case of Data Centres. Proceedings of the IEECB 2002. 2nd International Conference on Improving Electricity Efficiency in Commercial Buildings. Nice, France.

Energy per unit of floor area is not an adequate indicator for energy efficiency in high electric-load buildings. For data centres we propose to use a two-stage coefficient of energy efficiency $CEE = C1 * c2$, where $C1$ is a measure of the efficiency of the central infrastructure and $c2$ a measure of the energy efficiency of the equipment.

Anonymous. 2001. Model Data Center Energy Design Meeting. Austin Energy, Austin, TX, Feb 12-13. http://www.austinenergy.com/business/energy_design_meeting.htm

Anonymous. 2002a. 7 x 24 Update: Design & Construction - Issues and trends in mission critical infrastructure design, planning and maintenance.

<http://www.facilitiesnet.com/BOM/Jan02/jan02construction.shtml>. July 23, 2002.

<http://www.7x24exchange.org/>.

Anonymous. 2002b. Continuous Availability Review (CAR). The Uptime Institute: Computersite Engineering, Inc. <http://www.upsite.com/csepages/cseccar.html>. July 22, 2002.

Anonymous. 2002c. End-to-End Reliability Begins with the User's Definition of Success. The Uptime Institute. <http://www.upsite.com/TUIpages/editorials/endtoend.html>. July 22, 2002.

Anonymous. 2002d. Mechanical Systems Diagnostic Review (MSDR). The Uptime Institute: Computersite Engineering, Inc. <http://www.upsite.com/csepages/csemsdr.html>. July 22, 2002.

Anonymous. 2002e. Site Infrastructure Operations Review (SIOR). The Uptime Institute: Computersite Engineering, Inc. <http://www.upsite.com/csepages/cseior.html>. July 22, 2002.

Baer, D. B. Emerging Cooling Requirements & Systems in Telecommunications Spaces, Liebert Corporation.

During the last several years, power density trends, and consequently thermal density trends in telecommunications spaces have become topics of increasing interest. This paper identifies several of the underlying drivers of these trends, project possible outcomes, and assess the impact on cooling system design for these spaces.

Beck, F. 2001. Energy Smart Data Centers: Applying Energy Efficient Design And Technology To The Digital Information Sector. *Renewable Energy Policy Project (REPP)*: Washington, DC. (November 2001 REPP).

Both utilities and data center owners face challenges in meeting electricity demand loads with required levels of reliability. However, the bursting of the high-tech stock bubble in 2000 and the 2001 U.S. economic downturn has slowed expansion of data centers. This provides time and an opportunity to examine data center construction and operational practices with an eye toward reducing their energy demands through use of energy efficient technologies and energy smart design practices. As the economy recovers and the next data center rush approaches, best practices can reduce energy use while maintaining or even increasing data center reliability. Energy demands of data centers that support the digital information- and communications-based economy need not be as high as some predict. In fact, data center power demands could be reduced by 20 percent with minimal efficiency efforts, and by 50 percent with more aggressive efficiency measures.

Blount, H. E., H. Naah, and E. S. Johnson. 2001. Data Center and Carrier Hotel Real Estate: Refuting the Overcapacity Myth. Lehman Brothers: TELECOMMUNICATIONS, New York, June 7, 2001. <http://www.lehman.com>

An exclusive study examining supply and demand trends for data center and carrier hotel real estate in North America. Lehman Brothers and Cushman & Wakefield have completed the first in a regular series of proprietary studies on telecommunications real estate (TRE), including carrier hotels and data centers.

Bors, D. 2000. Data centers pose serious threat to energy supply. *Puget Sound Business Journal (Seattle)* - October 9, 2000. <http://seattle.bizjournals.com/seattle/stories/2000/10/09/focus5.html>

To cope with increasing energy demand from data centers, the author discussed feasibilities of two possible approaches: 1) energy industry approach by looking at

alternative energy supply; 2) construction industry approach by looking at data center energy efficiency. To get there, it is worth investigating five distinct components: (I) Co-generation of power. Presently, standby diesel generators are required to maintain the desired level of reliability at most data center sites, but their exhaust makes most of these generators unacceptable for long-term power generation; (II) Fuel cells offer the promise of very clean emissions and the reasonable possibility for use as standby power; (III) Increased efficiency in data center power distribution systems. There are two separate items that are major contributors to data center power distribution system inefficiencies. The first, power distribution units (PDUs), are available with optional internal transformers that use less energy than the present cadre of K-rated transformers. The second, uninterruptible power systems (UPSs), come in a range of efficiency ratings. If the use of high-efficiency PDUs and UPSs are combined, they offer the potential of a 6 percent saving; (IV) Increased efficiency in mechanical cooling systems. In order to ensure data center reliability, mechanical equipment is often selected as a large number of small, self-contained units, which offers opportunities to improve efficiencies; (V) Reductions in energy use by computer, network and storage equipment. Computer manufacturers can do their part by creating computers with greater computational power per watt. They have been doing this for years as a side effect of hardware improvements, and they can do even better if they make it a goal.

Brown, E., R. N. Elliott, and A. Shipley. 2001. Overview of Data Centers and Their Implications for Energy Demand. Washington, DC, American Council for an Energy Efficient Economy, Center for Energy & climate Solutions (CECS). September 2001.
<http://www.aceee.org/pdfs/datacenter.pdf.pdf>

The white paper discusses data center industry boom and energy efficiency opportunities and incentives in internet data centers. Emerging in the late 1990's, data centers are locations of concentrated Internet traffic requiring a high-degree of power reliability and a large amount of power relative to their square footage. Typically, power needs range from 10-40MW per building, and buildings are typically built in clusters around nodes in the Internet fiber-optic backbone. During the development boom in 1999 and 2000, projects averaged 6-9 months from site acquisition to operation, and planned operational life was 36 months to refit. Even high energy-prices were dwarfed by net daily profits of 1-2 million dollars per day for these buildings during the boom, creating little incentive for efficient use of energy.

Callsen, T. P. 2000. The Art of Estimating Loads. *Data Center* (Issue 2000.04).

This article discusses the typical Data Center layout. It includes floor plan analysis, HVAC requirements, and the electrical characteristics of the computer hardware typically found in a Data Center.

Calwell, C., and T. Reeder. 2002. Power Supplies: A Hidden Opportunity for Energy Savings (An NRDC Report). Natural Resources Defense Council, San Francisco, CA, May 22, 2002. <http://www.nrdc.org>

The article discusses the efficiency of power supplies which perform current conversion and are located inside of the electronic product (internal) or outside of the product (external). The study finds that most external models, often referred to as "wall-packs" or "bricks," use a very energy inefficient design called the linear power supply, with measured energy efficiencies ranging from 20 to 75%; that most internal power supply models use somewhat more efficient designs called switching or switch-mode power supplies; and that internal power supplies have energy efficiencies ranging from 50 to 90%, with wide variations in power use among similar products. Most homes have 5 to 10 devices that use external power supplies, such as cordless phones and answering machines. Internal power supplies are more prevalent in devices that have greater power requirements, typically more than 15 watts. Such devices include computers, televisions, office copiers, and stereo components. The paper points out that power supply

efficiency levels of 80 to 90% are readily achievable in most internal and external power supplies at modest incremental cost through improved integrated circuits and better designs.

Cratty, W., and W. Allen. 2001. Very High Availability (99.9999%) Combined Heat and Power for Mission Critical Applications. *Cinintel 2001*: 12. <http://www.surepowersystem.com>

Elliot, N. 2001. Overview of Data Centers and their implications for energy demand. Washington, DC, American Council for an Energy Efficient Economy. Jan 2001, revised June 10, 2001.

Feng, W., M. Warren, and E. Weigle. 2002. The Bladed Beowulf: A Cost-Effective Alternative to Traditional Beowulfs. *Cluster2002 Program*. <http://www-unix.mcs.anl.gov/cluster2002/schedule.html>; public.lanl.gov/feng/Bladed-Beowulf.pdf
Authors present a novel twist to the Beowulf cluster - the Bladed Beowulf. In contrast to traditional Beowulfs which typically use Intel or AMD processors, the Bladed Beowulf uses Transmeta processors in order to keep thermal power dissipation low and reliability and density high while still achieving comparable performance to Intel- and AMD-based clusters. Given the ever increasing complexity of traditional super-computers and Beowulf clusters; the issues of size, reliability, power consumption, and ease of administration and use will be "the" issues of this decade for high-performance computing. Bigger and faster machines are simply not good enough anymore. To illustrate, Authors present the results of performance benchmarks on the Bladed Beowulf and introduce two performance metrics that contribute to the total cost of ownership (TCO) of a computing system - performance/power and performance/space.

Frith, C. 2002. Internet Data Centers and the Infrastructure Require Environmental Design, Controls, and Monitoring. *Journal of the IEST* **45**(2002 Annual Edition): 45-52.
Internet Data Centers and the Infrastructure Require Environmental Design, Controls, and Monitoring. The author points out that specifications and standards need to be developed to achieve high performance for mission-critical internet applications.

Gilleskie, R. J. 2002. The Impact of Power Quality in the Telecommunications Industry. Palm Springs, CA, June 4. <http://www.energy2002.ee.doe.gov/Facilities.htm>
The workshop addresses the unique issues and special considerations necessary for improving the energy efficiency and reliability of high-tech data centers. This presentation addresses impacts of power quality including voltage sags, harmonics, and high-frequency grounding in telecommunication industry.

Grahame, T., and D. Kathan. 2001. Internet Fuels Shocking Load Requests. *Electrical World* **Vol. 215** (3): 25-27. http://www.platts.com/engineering/ew_back_issues.shtml
This article discusses the implications of the increase for power demand by the Internet's traffic growth on utility planning, operation, and financing.

Greenberg, D. 2001. Addendum to ER-01-15: A Primer on Harmonics. E-SOURCE, Boulder, Colorado, September 2001.
The electrical distribution systems of most commercial and industrial facilities were not designed to operate with an abundance of harmonics-producing loads. In fact, it is only within recent years that such loads have become widespread enough for industry to take notice and to begin to develop strategies to address the problems that harmonics can create. By 1992, concern about the issue had grown sufficiently that the Institute for Electrical and Electronic Engineers (IEEE) developed and published its standard 519, "IEEE Recommended Practices and Requirements for Harmonic Control in Electrical Power Systems," which established an approach for setting limits on the harmonic voltage distortion on the utility power system and on the harmonic currents created individual power consumers. Since that time, the electronic loads that give rise to harmonic currents have grown dramatically and are projected to continue growing for the

foreseeable future. This being the case, there is and will continue to be a market for technological solutions to the problems that harmonics can cause.

Gross, P. 2002. Needed: New Metrics. *Energy User News*.

http://www.energyusernews.com/eun/cda/articleinformation/features/bnp__features__item/0,2584,82741,00.html

Gruener, J. 2000. Building High-Performance Data Centers. *Dell Magazines - Dell Power Solutions* (Issue 3 "Building Your Internet Data Center").

http://www.dell.com/us/en/esg/topics/power_ps3q00_1_power.htm;

http://www.dell.com/us/en/esg/topics/power_ps3q00-giganet.htm

The introduction of Microsoft SQL Server 2000 is a milestone in the race to build the next generation of Internet data centers. These new data centers are made up of tiers of servers, now commonly referred to as server farms, which generally are divided into client services servers (Web servers), application/business logic servers, and data servers supporting multiple instances of databases such as SQL Server 2000.

Hellmann, M. 2002. Consultants Face Difficult New Questions in Evolving Data Center Design. *Energy User News*.

http://www.energyusernews.com/CDA/ArticleInformation/features/BNP__Features__Item/0,2584,70610,00.html

While few data center design projects are alike, there are always the twin challenges of "power and fiber." And sometimes, even local politics and human factors. The paper suggested that the consultant should be brought in as soon as a business case is established so criteria can be established and a concept can be developed, priced, and compared to the business case. A planning is necessary before moving on to site selection and refine the concept and again test the business case.

Howe, B., A. Mansoor, and A. Maitra. 2001. Power Quality Guidelines for Energy Efficient Device Application - Guidebook for California Energy Commission (CEC). Final Report to B. Banerjee, California Energy Commission (CEC).

Energy efficiency and conservation are crucial for a balanced energy policy for the Nation in general and the State of California. Widespread adaptation of energy efficient technologies such as energy efficient motors, adjustable speed drives, improved lighting technologies will be the key in achieving self sufficiency and a balanced energy policy that takes into account both supply side and demand side measures. In order to achieve the full benefit of energy efficient technologies, these must be applied intelligently, and with clear recognition of the impacts some of these technologies may have on power quality and reliability. Any impediment to the application of these energy efficient technologies by the customers is not desirable for the overall benefit to energy users in California. With that in mind EPRI and CEC has worked to develop this guidebook to promote customer adaptation of energy efficient technologies by focusing on three distinct objectives. 1) Minimize any undesirable power quality impacts of energy-saving technologies; 2) Understand the energy savings potential of power quality-related technologies. These include: Surge Protective Devices (SPDs) or Transient Voltage Surge Suppressors (TVSS), Harmonic Filters, Power Factor Correction Capacitors, Electronic Soft Starters for Motors; and 3) How to evaluate "black box" technologies

Intel. 2002. Planning and Building a Data Center - Meeting the e-Business Challenge. Intel Corp. http://www.intel.com/network/idc/doc_library/white_papers/data_center/. Aug 01, 2002.

The paper discusses the keys to success of Internet Service Providers (ISPs) that include 1) Achieve the economies of scale necessary to support a low price business model; 2) Offer added value, typically in the form of specialized services such as applications hosting to justify a premium price. This document provides a high-level overview of the requirements for successfully establishing and operating an Internet

data center in today's marketplace. It offers some of the key steps that need to be taken, including project definition, prerequisites and planning. In order to construct a data center that can meet the challenges of the new market, there are three basic areas of data center definition and development: 1) Facilities: including building, security, power, air-conditioning and room for growth; 2) Internet connectivity: performance, availability and scalability; 3) Value-added services and the resources to support their delivery: service levels, technical skills and business processes. The aim is to provide customers with the physical environment, server hardware, network connectivity and technical skills necessary to keep Internet business up and running 24 hours a day, seven days a week. The ability to scale is essential, allowing businesses to upgrade easily by adding bandwidth or server capacity on demand.

Koplin, E. 2000. Finding Holes In The Data Center Envelope. *Engineered Systems* (September 2000).

http://www.esmagazine.com/CDA/ArticleInformation/features/BNP__Features__Item/0,2503,8720,00.html

The paper addresses importance of environmental control in data center facilities. Maintaining data center availability requires absolutely reliable infrastructure. A significant amount of this is devoted solely to maintaining stable environmental parameters. And only constant, thorough regulation and testing of these parameters ensures the integrity of the data center "envelope."

Mandel, S. 2001. Rooms that consume - Internet hotels and other data centers inhale electricity. *Electric Perspectives* Vol. 26 (No.3).

http://www.eei.org/ep/editorial/Apr_01/0401ROOM.htm

The article estimated that the amount of this data center space in the United States nearly doubled in 2000, totaling between 19 million and 25 million square feet by year-end, according to investment analysts. They say they expect another 10 million to 20 million square feet of new space to be added in 2001. Developers are asking electric utilities to supply the buildings with 100-200 watts of electricity per square foot. Since these data centers are new to the economy, there is little historical data on which to base estimates of electricity use for a facility. In addition, the dot.com world makes it difficult for the developer to say confidently how much electricity one of these internet hotels will use. Source One estimates that tens of billions of dollars worth of electric infrastructure improvements will be needed for data centers over the next few years and that they will consume billions of dollars more worth of electricity. The energy costs are as high or higher than the actual lease costs. Indeed, 50-60 percent of the cost of building a data center is for the power, including batteries, backup generators, and air-conditioning, as well as the cost for utility construction.

Mitchell-Jackson, J. 2001. Energy Needs in an Internet Economy: A Closer Look at Data Centers, July, 2001.

This study explains why most estimates of power used by data centers are significantly too high, and gives measured power use data for five such facilities. Total power use for the computer room area of these data centers is no more than 40 W/square foot, including all auxiliary power use and cooling energy. There are two draft journal articles from this work, one focusing on the detailed power use of the data center we've examined in most detail, and the other presenting the aggregate electricity use associated with hosting-type data centers in the U.S.

Mitchell-Jackson, J., J. G. Koomey, B. Nordman, and M. Blazek. 2001. Data Center Power Requirements: Measurements From Silicon Valley. *Energy—the International Journal (Under review)*. <http://enduse.lbl.gov/Projects/InfoTech.html>

Current estimates of data center power requirements are greatly overstated because they are based on criteria that incorporate oversized, redundant systems, and several safety factors. Furthermore, most estimates assume that data centers are filled to capacity. For the most part, these numbers are unsubstantiated. Although there are many estimates of the amount of electricity consumed by data centers, until this study, there were no publicly available measurements of power use. This paper examines some of

the reasons why power requirements at data centers are overstated and adds actual measurements and the analysis of real-world data to the debate over how much energy these facilities use.

Nordham, Reiss, and Stein. 2001. Delivering Energy Services to Internet Hotels and Other High Density Electronic Loads, Part I: Structure of the HiDEL Industry. Platts Research and Consulting, Boulder, CO.

Patel, C. D., C. E. Bash, C. Belady, L. Stahl, and D. Sullivan. 2001. Computational Fluid Dynamics Modeling of High Compute Density Data Centers to Assure System Inlet Air Specifications. Reprinted from the proceedings of the Pacific Rim ASME International Electronic Packaging Technical Conference and Exhibition (IPACK 2001), © 2001, ASME.

Due to high heat loads, designing the air conditioning system in a data center using simple energy balance is no longer adequate. Data center design cannot rely on intuitive design of air distribution. It is necessary to model the airflow and temperature distribution in a data center. This paper presents a computational fluid dynamics model of a prototype data center to make the case for such modeling.

Patel, C. D., R. Sharma, C. E. Bash, and A. Beitelmal. 2002. Thermal Considerations in Cooling Large Scale High Compute Density Data Centers. 8th ITherm Conference. San Diego CA.

A high compute density data center of today is characterized as one consisting of thousands of racks each with multiple computing units. The computing units include multiple microprocessors, each dissipating approximately 250 W of power. The heat dissipation from a rack containing such computing units exceeds 10 KW. Today's data center, with 1000 racks, over 30,000 square feet, requires 10 MW of power for the computing infrastructure. A 100,000 square foot data center of tomorrow will require 50 MW of power for the computing infrastructure. Energy required to dissipate this heat will be an additional 20 MW. A hundred thousand square foot planetary scale data center, with five thousand 10 KW racks, would cost ~\$44 million per year (@ \$100/MWh) just to power the servers & \$18 million per year to power the cooling infrastructure for the data center. Cooling design considerations by virtue of proper layout of racks can yield substantial savings in energy. This paper shows an overview of a data center cooling design and presents the results of a case study where layout change was made by virtue of numerical modeling to avail efficient use of air conditioning resources.

PG&E. 2001. Data Center Energy Characterization Study. Pacific Gas and Electric Company (subcontractor: Rumsey Engineers), San Francisco, Feb. 2001.

Rumsey Engineers, Inc. and PG&E have teamed up to conduct an energy study as part of PG&E's Data Center Energy Characterization Study. This study will allow PG&E and designers to make better decisions about the design and construction of data centers in the near future. Three data centers in the PG&E service territory have been analyzed during December 2000 and January 2001, with the particular aim of determining the end-use of electricity. The electricity use at each facility was monitored for a week each. At the end of the report are a set of definitions, which explain the terms used and the components in making each calculation. The three data centers provide co-location service, which is an unmanaged service that provides rack space and network connectivity via a high capacity backbone. About half or more of the electricity goes to powering the data center floor, and 25 to 34 percent of the electricity goes to the heating, air conditioning and ventilation equipment. The HVAC equipment uses a significant amount of power and is where energy efficiency improvements can be made. All three facilities use computer room air conditioning (CRAC) units, which are stand-alone units that create their own refrigeration and circulate air. A central, water-cooled chilled water system with air handlers and economizers can provide similar services with roughly a 50% reduction in cooling energy consumption. Energy density of the three buildings had an average of 35 W/sf. The cooling equipment energy density for the data center floor alone averaged at 17 W/sf for the three facilities. The average designed energy density of the three data centers' server loads was 63 W/sf, while the measured energy density was 34 W/sf. An extrapolated value was also calculated to determine what the server load energy density would be when fully occupied. The average

extrapolated energy density was 45 W/sf. Air movement efficiency varies from 23 to 64 percent between the three facilities. Cooling load density varies from 9 to 70 percent between the three facilities.

Planet-TECH. 2002. Technical and Market Assessment for Premium Power in Haverhill. Planet-TECH Associates for The Massachusetts Technology Collaborative, www.mtpc.org, Westborough, MA 01581-3340, Revision: February 20, 2002.

http://www.mtpc.org/cluster/Haverhill_Report.pdf ; <http://www.planet-tech.com/content.htm?cid=2445>

This study is pursued under contract to the Massachusetts Technology Collaborative, in response to a request for a "Technical and Market Assessment". It seeks to determine if the provisioning of "premium power" suitable for data-intensive industries will improve the marketability of a Historic District mill building in Haverhill. It is concluded that such provisioning does improve the marketability, however, not to a degree that is viable at this time. Other avenues for energy innovation are considered and recommendations for next steps are made.

RMI. 2003. Draft Final Report on Energy Efficient Data Centers - A Rocky Mountain Institute Design Charrette. Hayes Mansion Conference Center, San Jose, California.

Rapid growth of "mission critical" server-farm and fiber-optic-node data centers has presented energy service providers with urgent issues. Resulting costs have broad financial and societal implications. While recent economic trends have severely curtailed projected growth, the underlying business remains vital. This charrette, hosted and coordinated by the Rocky Mountain Institute (RMI- www.rmi.org), assembled over 75 leading stakeholders in the data center market for three days of brainstorming and sharing their expertise and information. The charrette participants identified many issues associated with efficient data center design. Many ideas for needed research were presented.

Robertson, C., and J. Romm. 2002. Data Centers, Power, and Pollution Prevention - Design for Business and Environmental Advantage. The Center for Energy and Climate Solutions; A Division of The Global Environment and Technology Foundation, June 2002.

<http://www.cool-companies.org>; <http://www.getf.org>

Computers and other electronic equipment will crash at the slightest disruption or fluctuation in their supply of electricity. The power system was not designed for these sensitive electronic loads and is inherently unable to meet the technical requirements of the information economy. For data centers, which play a central role in the information economy, crashing computers cause potentially catastrophic financial losses. The same voltage sag that causes the lights to dim briefly can cause a data center to go off-line, losing large sums of money, for many hours. Data center owners and their power providers must therefore solve several related technical and economic electric power problems. These are: 1) How to assure high-availability (24x7) power supply with a very low probability of failure; 2) How to assure practically perfect power quality; and 3) How to manage risk while minimizing capital and operating expenses

Roth, K. W., Fred Goldstein, and J. Kleinman. 2002. Energy consumption by office and telecommunications equipment in commercial buildings, Volume I: Energy Consumption Baseline. Arthur D. Little (ADL), Inc., 72895-00, Cambridge, MA, January 2002.

ADL carried out a "bottom-up" study to quantify the annual electricity consumption (AEC) of more than thirty (30) types of non-residential office and telecommunications equipment. A preliminary AEC estimate for all equipment types identified eight key equipment categories that received significantly more detailed studied and accounted for almost 90% of the total preliminary AEC. The Key Equipment Categories include: *Computer Monitors and Displays, Personal Computers, Server Computers, Copy Machines, Computer Network Equipment, Telephone Network Equipment, Printers, Uninterruptible Power Supplies (UPSs)*. The literature review did not uncover any prior comprehensive studies of telephone network electricity consumption or uninterruptible power supply (UPS) electricity consumption. The AEC analyses found that the office and telecommunications equipment consumed 97-TWh of electricity in 2000. The report concludes that commercial sector office equipment electricity use in the U.S. is about 3% of all

electric power use. The ADL work also creates scenarios of future electricity use for office equipment, including the energy used by telecommunications equipment.

Shields, H. and C. Weschler, 1998. Are Indoor Pollutants Threatening the Reliability of Your Electronic Equipment? Heating/Piping/Air Conditioning Magazine. May.

Stein, Jay. 2002. More Efficient Technology Will Ease the Way for Future Data Centers. Proceedings 2002 ACEEE Summer Study on Energy Efficiency in Buildings.

Sullivan, R. F. 2002. Alternating Cold and Hot Aisles Provides More Reliable Cooling for Server Farms. The Uptime Institute. <http://www.uptimeinstitute.org/tuiaisles.html>
The creation of "server farms" comprising hundreds of individual file servers has become quite commonplace in the new e-commerce economy, while other businesses spawn farms by moving equipment previously in closets or under desktops into a centralized data center environment. However, many of these farms are hastily planned and implemented as the needed equipment must be quickly installed on a rush schedule. The typical result is a somewhat haphazard layout on the raised floor that can have disastrous consequences due to environmental temperature disparities. Unfortunately, this lack of floor-layout planning is not apparent until after serious reliability problems have already occurred.

The Uptime Institute. 2000. Heat-Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment. The Uptime Institute, Version 1.0., <http://www.uptimeinstitute.org/heatdensity.html>
This white paper provides data and best available insights regarding historical and projected trends in power consumption and the resulting heat dissipation in computer and data processing systems (servers and workstations), storage systems (DASD and tape), and central office-type telecommunications equipment. The topics address the special needs of Information Technology professionals, technology space and data center owners, facilities planners, architects, and engineers.

Thompson, C. S. 2002. Integrated Data Center Design in the New Millennium. *Energy User News*.
http://www.energyusernews.com/CDA/ArticleInformation/features/BNP__Features__Item/0,2584,70578,00.html
Data center design requires planning ahead and estimating future electrical needs. Designers must accurately predict space and energy requirements, plus cooling needs for new generations of equipment. Importance of data center reliability is discussed.

Uptime Institute, 2000. Heat Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment. Santa Fe, NM.

Wood, L. 2002. Cutting Edge Server Farms - The blade server debate. [newarchitectmag.com](http://www.newarchitectmag.com/documents/s=2412/na0702f/index.html).
<http://www.newarchitectmag.com/documents/s=2412/na0702f/index.html>. July 23, 2002.
A blade is the industry term for a server that fits on a single circuit board, including CPU, memory, and perhaps a local hard disk. Multiple blades are plugged into a chassis, where each blade shares a common power supply, cooling system, and communications back plane. Multiple chassis can then be stacked into racks. By comparison, the conventional approach for rack-mounted servers involves only one server per chassis. A chassis cannot be smaller than one vertical rack unit (1U, or about 1.75 inches high). This limits you to 42 to 48 servers in a standard seven-foot rack. A typical blade chassis is much higher than 1U, but several can still be stacked in a rack, allowing upwards of 300 servers per rack, depending on the vendor and configuration. This compact design offers compelling

advantages to anyone operating a high-density server farm where space is at a premium. Indeed, blades are the "next big thing" in servers, and it's probable that any given administrator will have to decide whether to adopt them in the near future.

Uptime Institute, 2000. Heat Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment. Santa Fe, NM.