Design Recommendations for

High-Performance

Data Centers

Report of the

Integrated Design Charrette

Conducted

2–5 February 2003

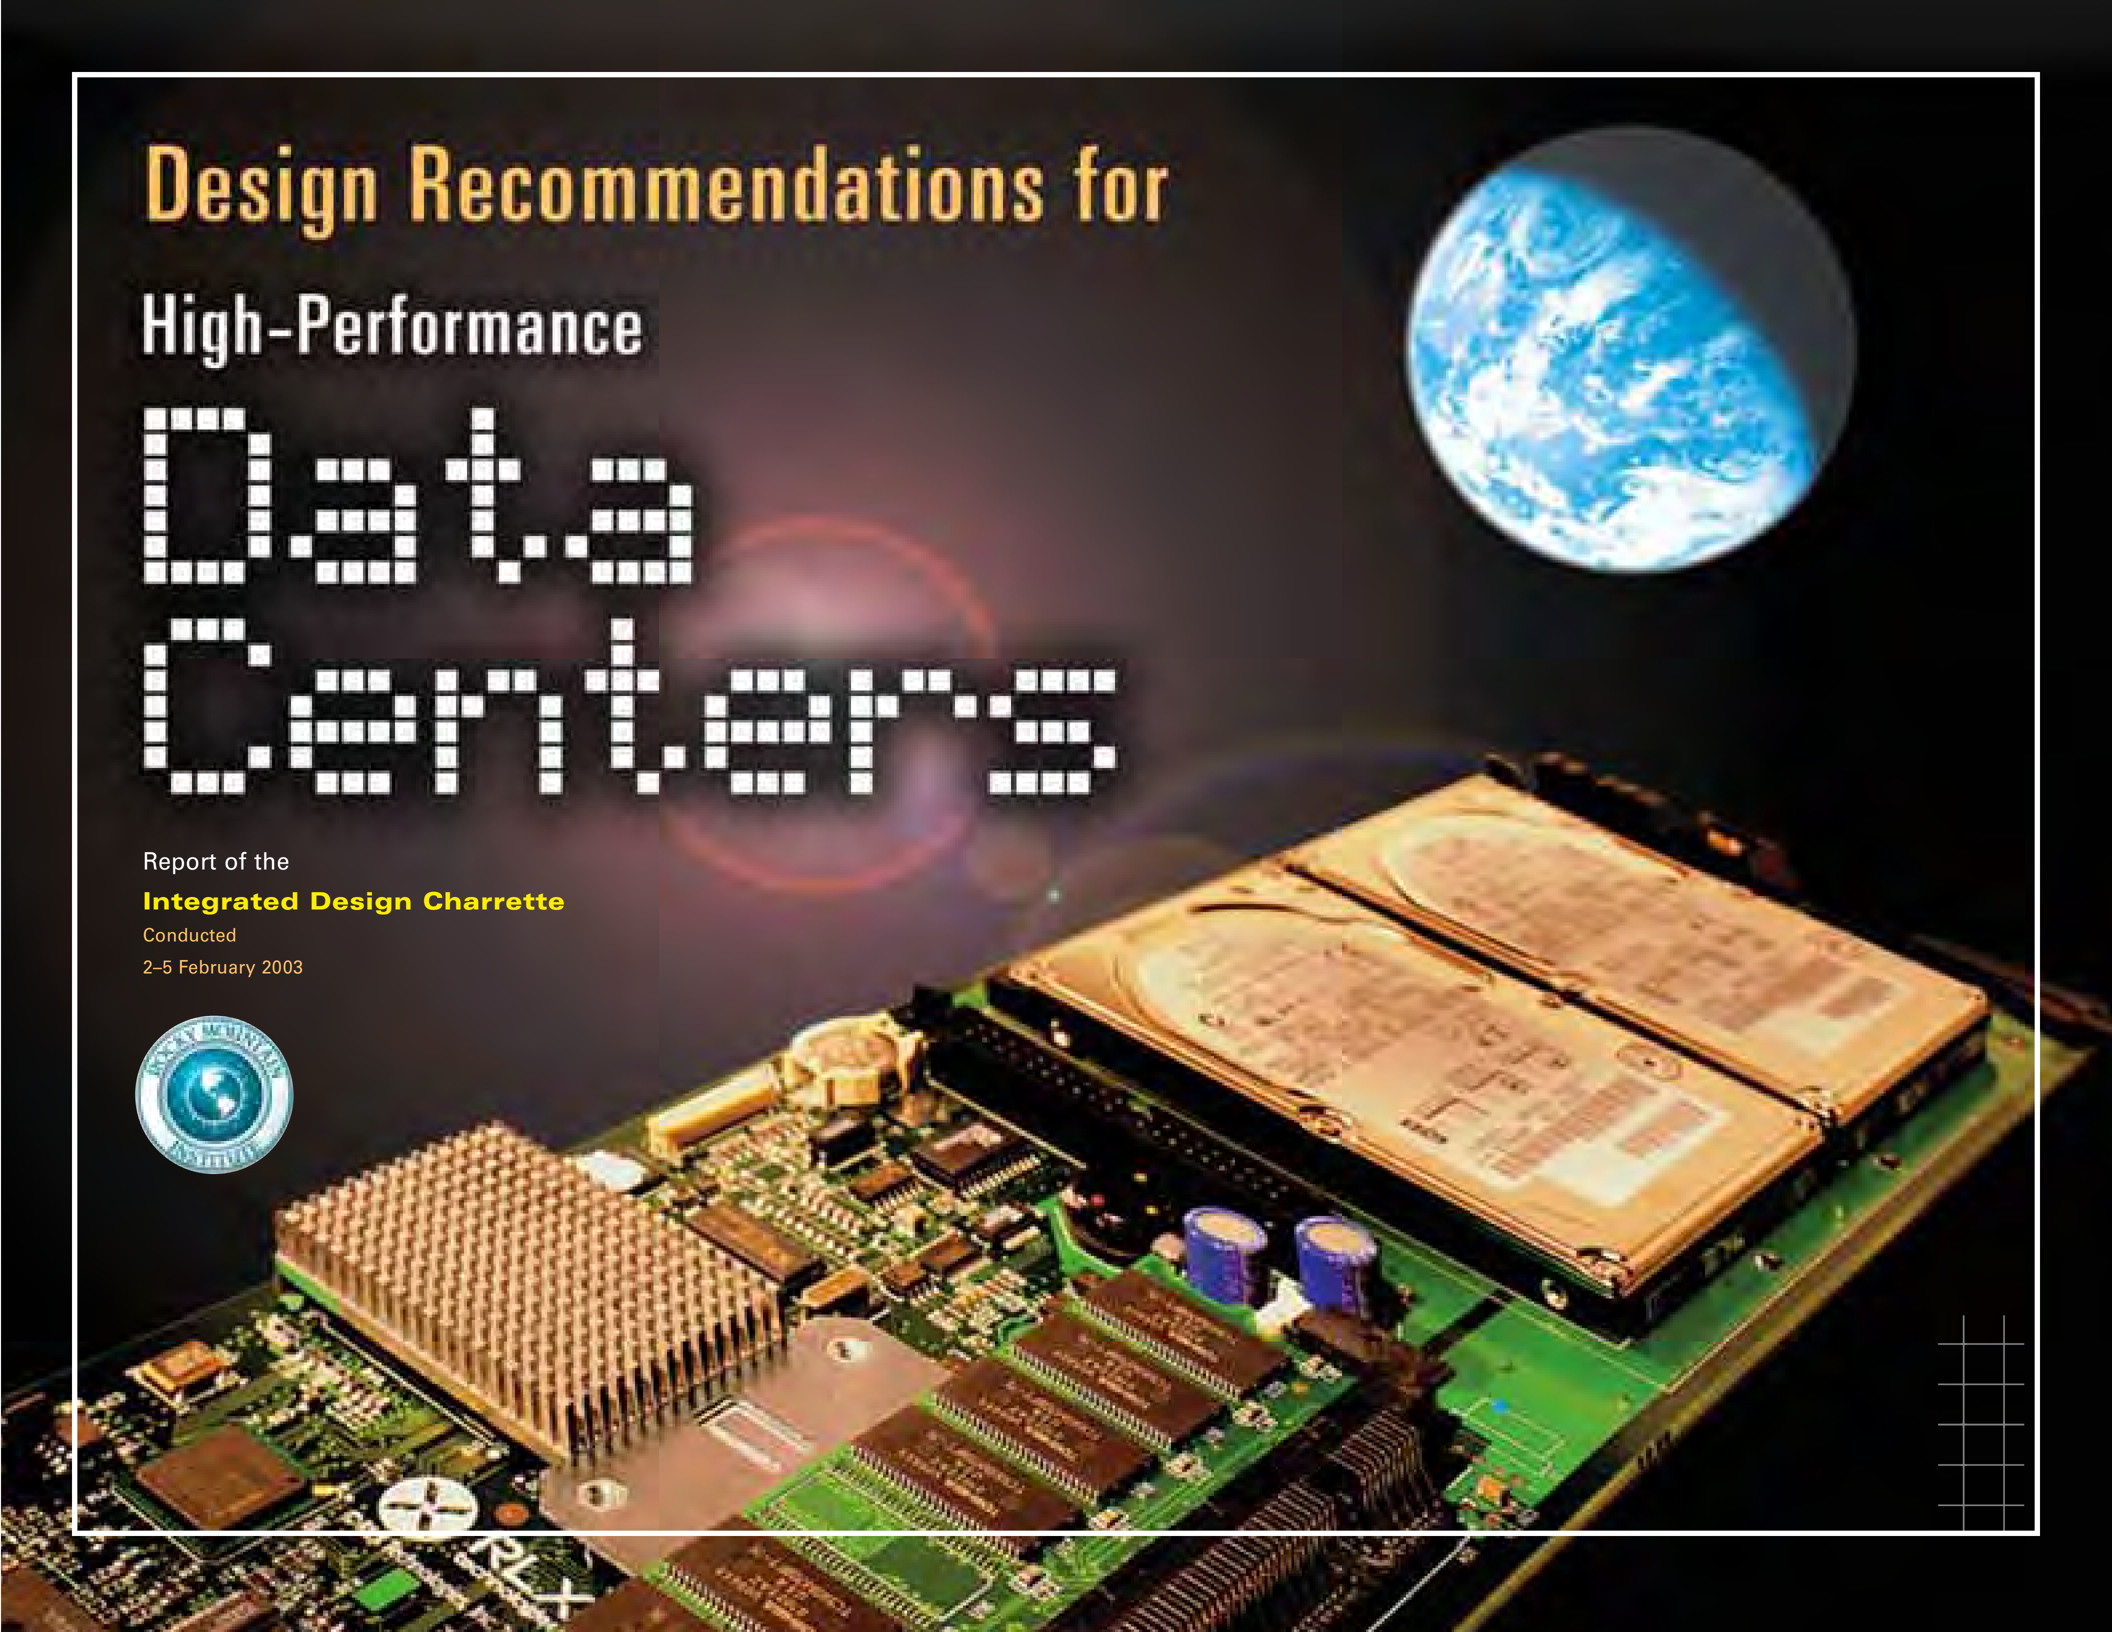

This document may be reprinted with proper credit
to Rocky Mountain Institute (RMI).
Please visit www.rmi.org/sitepages/pid626.php
for the report, addenda, and ongoing corrections and updates.

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.

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

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

Table of Contents

Design Recommendations for
High-Performance
Data Centers

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.

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)

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 Westbrock, 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

Introduction

> "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*

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

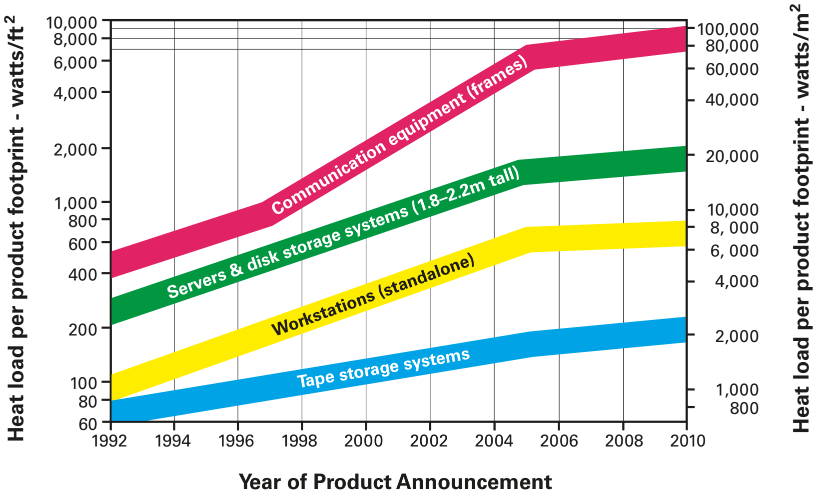

Source: Uptime Institute; www.uptime.com.

---

[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, Arrenhius' 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. Feng[4] 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 cm^3, not cm^2), 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, ehw_@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 Charrette[4]

On 2–5 February 2003, Rocky Mountain Institute (RMI—see www.rmi.org and p. 85) convened an innovative whole-system design charrette[5] 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

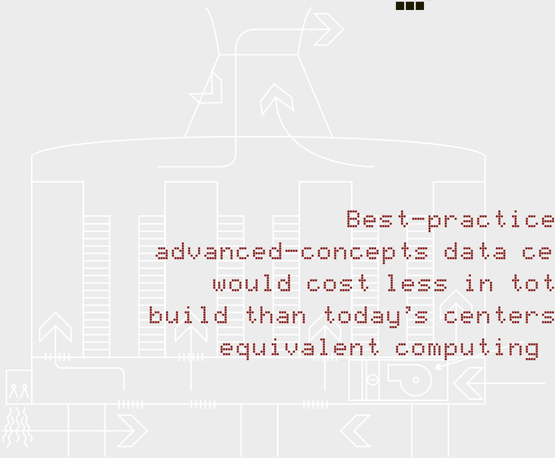

Best-practices and advanced-concepts data centers would cost less in total to build than today's centers with equivalent computing power

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

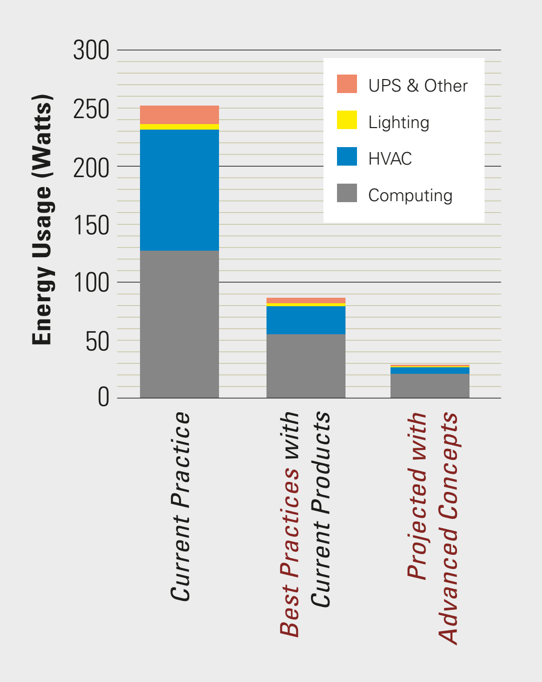

Source: Malcolm Lewis

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 Lewis[1] 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 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

---

[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.

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

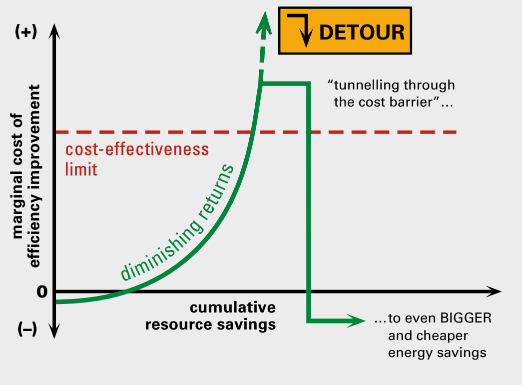

Source: Rocky Mountain Institute

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.

---

[2] Loosely quoted from *Natural Capitalism*, pp. 121–22.

Today's typical 1-U[3] rack-mounted server uses approximately 128 watts. The conceptual "Hyperserver" 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.

Computing End-Use Energy

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 70–90ºF—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 euthalpy-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 *and* 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.

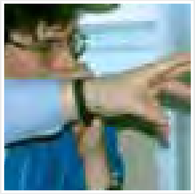

Operations

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 + 10^1× 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.

10^1: "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
- *Mmmmm!*

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/W[4]—more than for solar kWh
- Low-V (≤1 V), high-*k*, voltage islands, VLIW, SiGe, Cu, Si/ins; RLX now gets ~5–8× saving
- 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/m^2 for telco on a DC bus as for internet applications on an AC bus, perhaps partly due to that difference[5]

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,…

Figure S.10: PAX fan

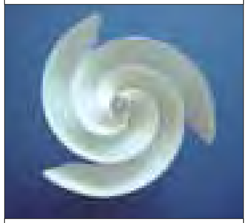

Source: PAX Scientific

- Water cooling? (could be dry-chip, plug-and-play; 3,467× heat cap/vol + 10^1× 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

The morsels, scraps, and broth

- Building envelope, α, ε, shading, elevators
- Lighting (~1–3 W/m^2 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±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 ≥3×? (*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

---

[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.

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). Comprehensible 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.

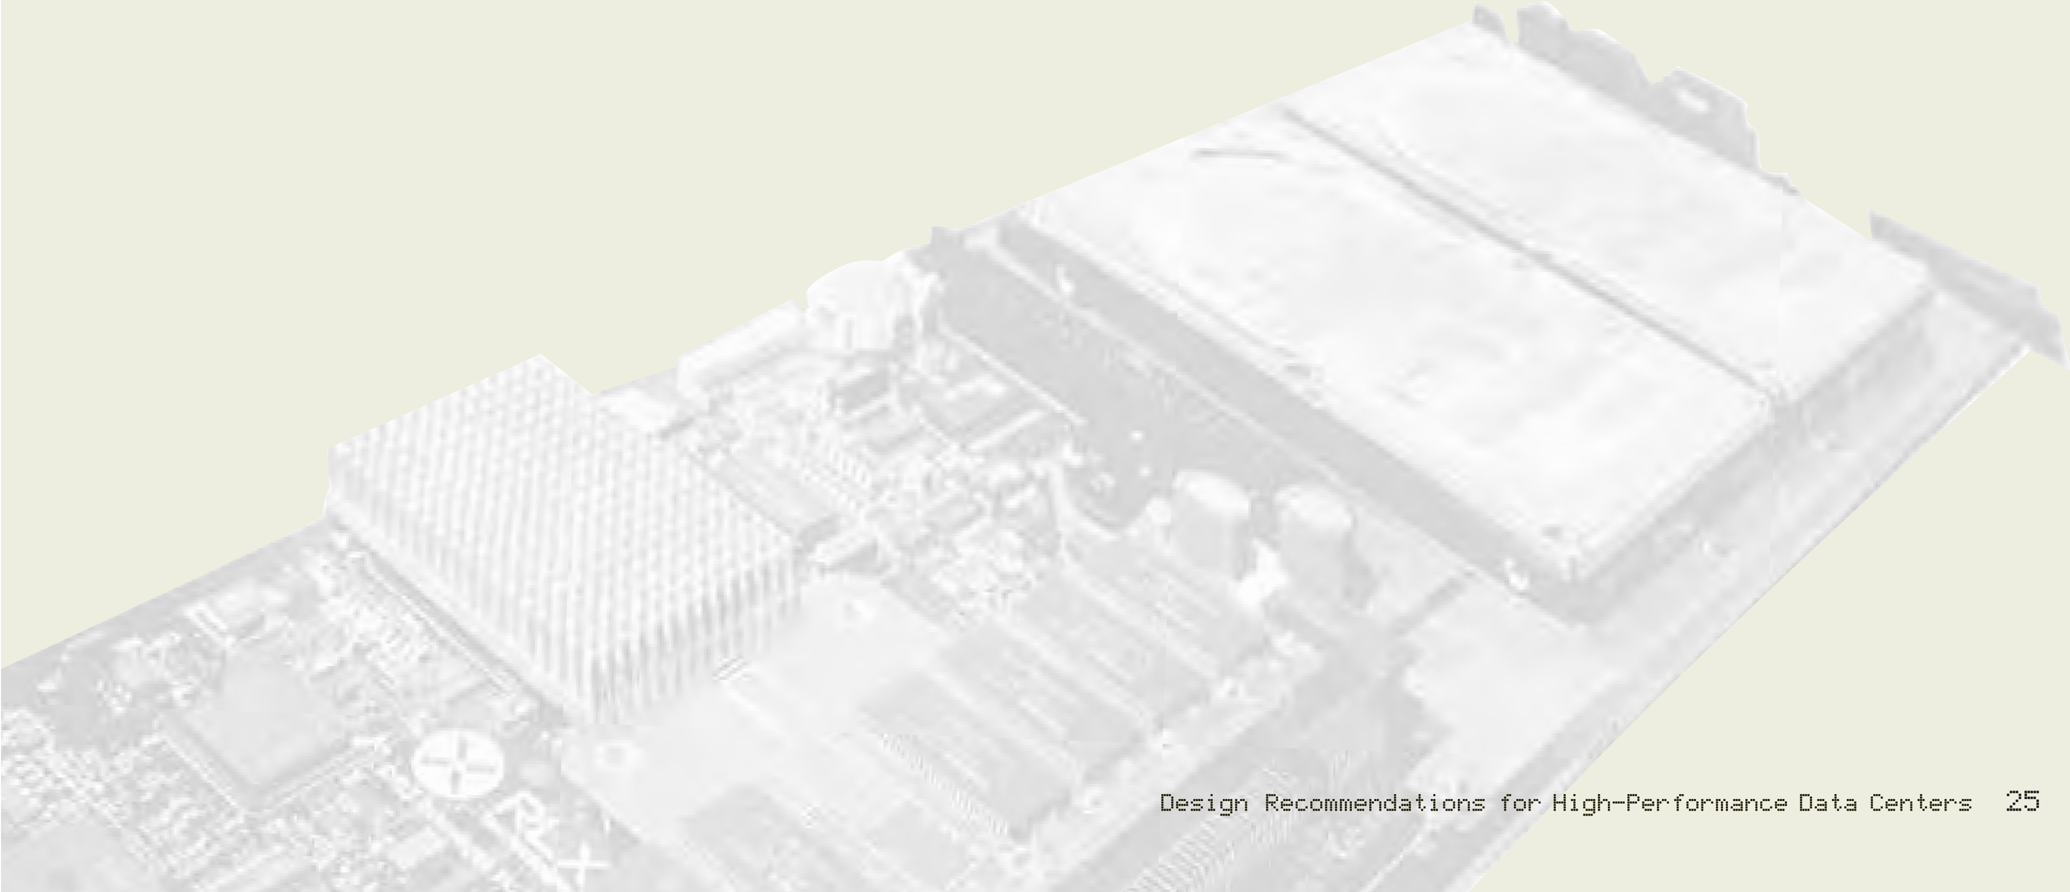

Recommendation	Available Now	In Development/Future	Equipment Designers/Manufacturers	Engineers/Architects	Building Owners/Developers	Real Estate Brokers	Policy/Industry Assoc. Trades	Information Technology People	Comments
Part 1 Native Loads (CPUs, Servers, Software)									
1.1 *New design paradigm: "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									*A 35–95 percent efficiency improvement can be achieved with a 40–400 percent/y ROI—with no decrease in reliability*
4.A Low Energy Ventilation & Cooling									
4.A.1 Increase the temperature range of cooling air	X		X	X				X	*Systems that boost air temperature differentials increase cooling system efficiency*
4.A.2 Manage airflow to reduce energy required for cooling and ventilation	X		X	X	X		X	X	
4.A.3 Minimize air side static pressure	X			X					*Reduce system resistances by making detailed improvements in dynamic flow paths and efficiencies*
4.A.4 Maximize use of free cooling	X			X					*Bringing in cold-climate outside air instead of cooling return air from inside the data center reduces cooling loads*
4.A.5 Natural ventilation	X			X	X	X	X		*Natural ventilation is suitable for data centers located in dry climates—either hot or cold*
4.A.6 Demand-controlled ventilation	X			X	X				
4.A.7 Additional ideas	X		X	X	X		X	X	*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*
4.A.8 Wish list for manufacturers		X	X				X	X	*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*

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		

Design Recommendations for High-Performance Data Centers

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 Loads[1] (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:

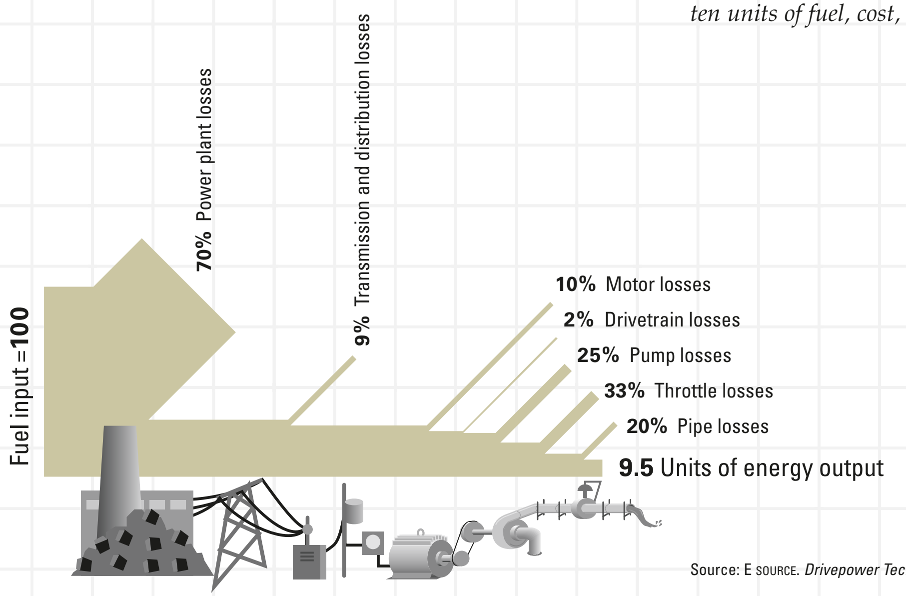

Source: E SOURCE. *Drivepower Technology Atlas* (www.esource.com).

> *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]

---

[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**	

Figure 1.1.2: Current practice

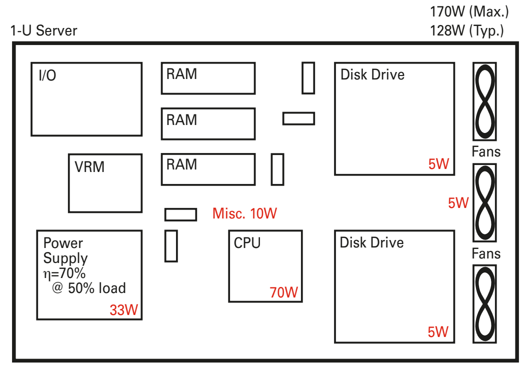

Of the 14W, the power converter(s) on the Hyperserver accounts for 3W. Although the fraction is high, in absolute terms the 3W are fairly easy to handle locally. The efficiency of conversion from AC to DC determines the heat generated in that process, based on the resistance of the rectifier and the 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

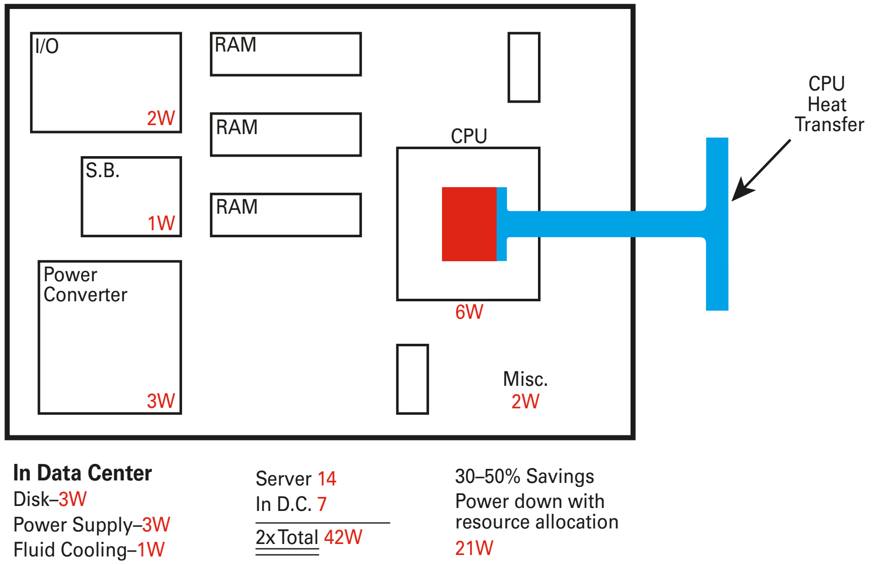

Figure 1.1.5: The Intel Itanium server

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.

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.

---

[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.

Recommendation

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).

Figure 1.4: Moore's Law for power dissipation

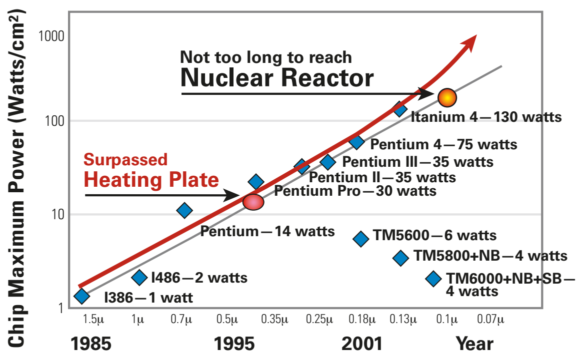

Source: Chris Hipp

Recommendation

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 "FLOPS[8]/cycle," rather than clock speed.

---

[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.

1.5 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.

Photo courtesy Chris Hipp

---

[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."

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 resources[10]

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.

Figure 2.1.1: Computer power supply efficiencies

Efficiency (%) as a function of load (%) of a modern power supply for servers (maximal load = 400 W)

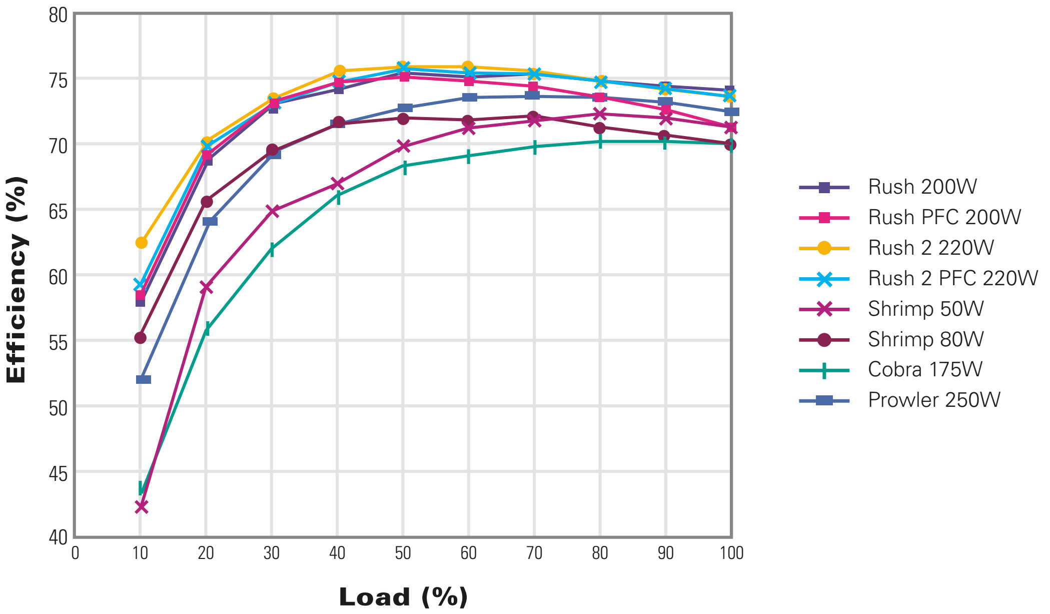

Source: NRDC report by Chris Calwell and Travis Reeder, 22 May 2002.

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 1W × 8766 hours/year x $0.07/kWh × 0.001 kW/W × 3-year lifetime × 2 W-input/W-delivered = $4/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.

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.

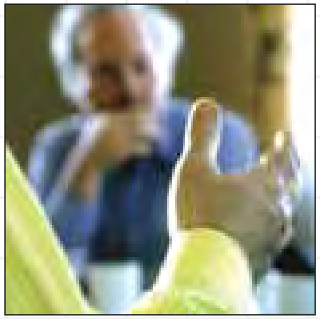

Size, Weight, and Cost

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 traces[4] 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 go 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 times[1] 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 concept[3] (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/ft^3 = 62.4 Btu/ft^3-F°.
The heat capacity of air is 0.24 Btu/lbm-F° * 0.075 lbm/ft^3 = 0.018 Btu/ft^3-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/VSWS188.pdf.)

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.

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.

Panel Comments on Cooling

Dick Bourne: I just did a little calculation. If you're willing to accept that a seven-[ºF-]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.

Recommendation

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).

Recommendation

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.
- 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 3.3.1: Near-term solution

Hybrid: Cools processor with fluid remainder air-cooled

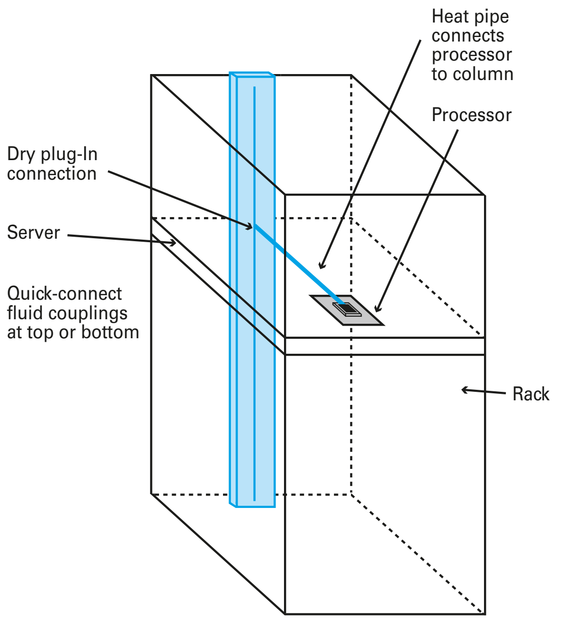

Figure 3.3.2: Near-term airless

- Works with pizza boxes or blades
- Horizontal transition accessory for blades
- As in hybrid—fluid column connects at top or bottom

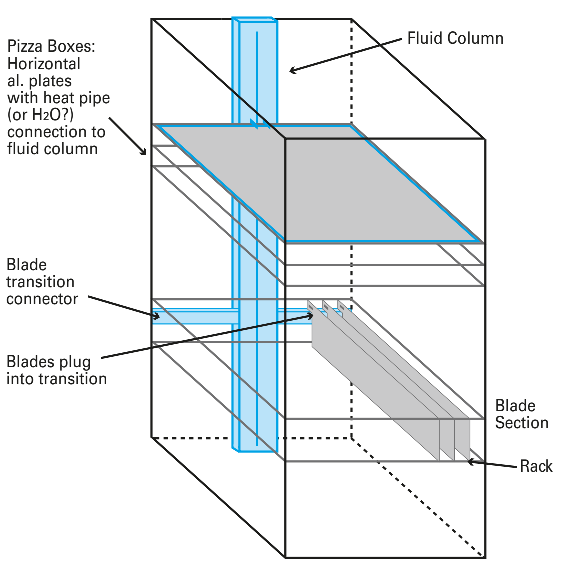

3.4 Conductive thermal path to liquid for future

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.

Figure 3.4: Long-term blade cooling
Airless: conductive and/or radiant "powerless": modular/separate from racks

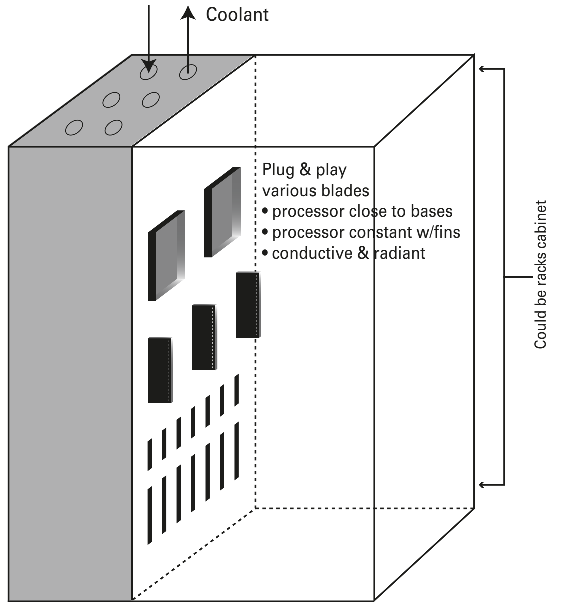

The recommended long-term solution, shown in Figure 3.4, is based on a "cool wall" located near or on the back wall of the rack. "Hot spots" (*e.g.*, processors) would be located close to or against that wall. Whatever heat conveyance system is then used must carry heat only a short distance. If blade-type 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.

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.
- 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.

Photo courtesy Chris Hipp

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

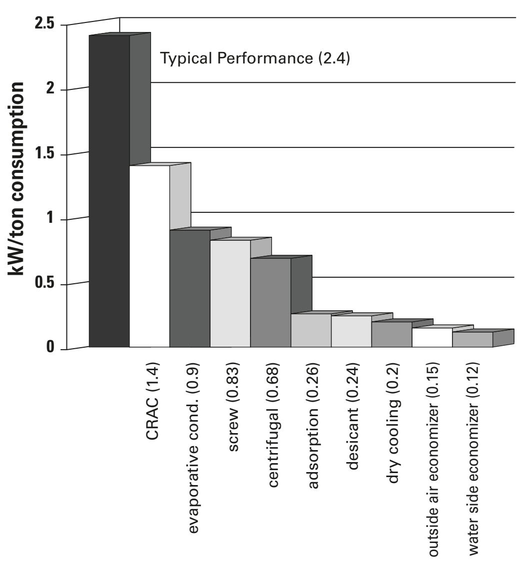

Figure 4.1.2: Cooling cost savings from more efficient conventional cooling systems

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, economizer[1] 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.

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 55–75ºF range to 70–90ºF—thereby improving efficiency and increasing the opportunities for economizer operations. Improving airflow through the equipment racks

Figure 4a.2.1

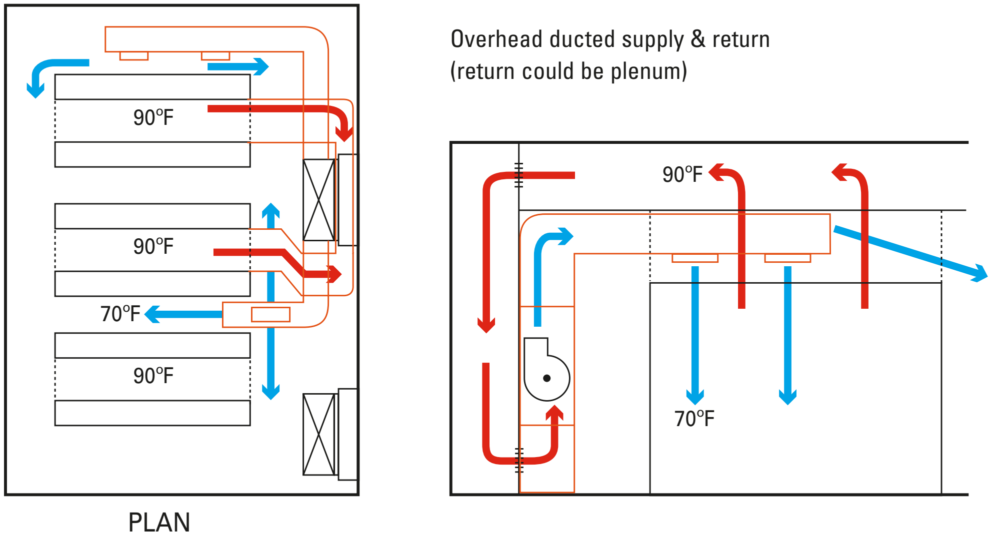

---

[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.

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

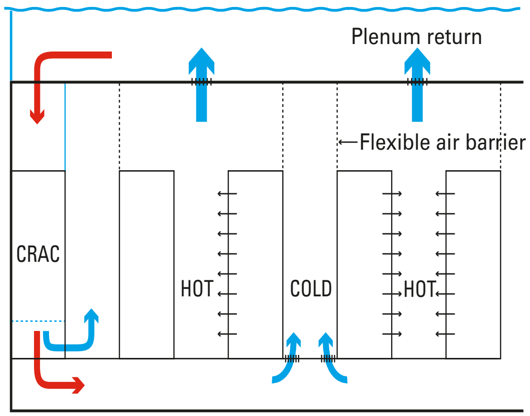

Alternating hot and cold aisles: Figures 4a.2.1–4 illustrate various methods for establishing alternate hot and cold aisles that provide the temperature differentials necessary to draw heat efficiently through the computer equipment to the return air plenum. This arrangement also avoids mixing hot and 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 underfloor 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 barrier 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-cool 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

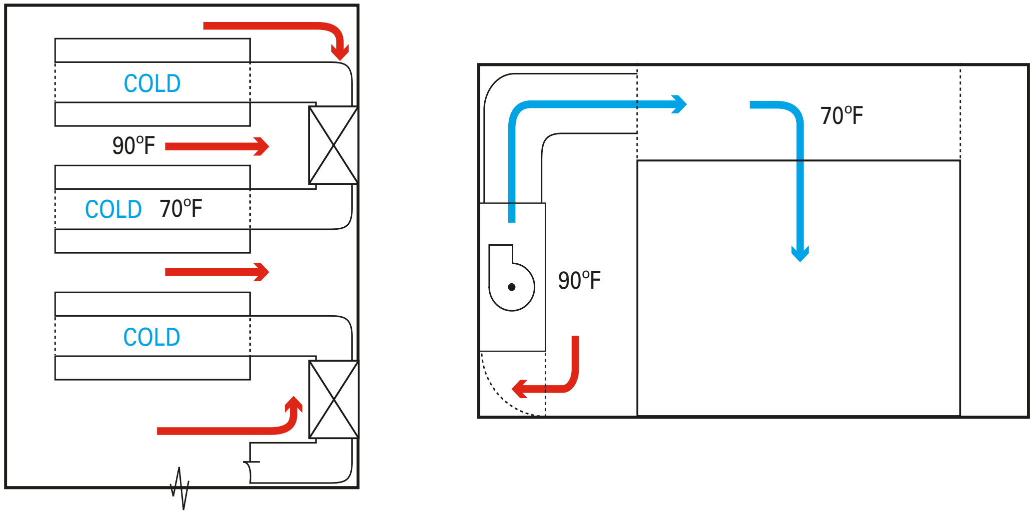

Figure 4.a.2.4

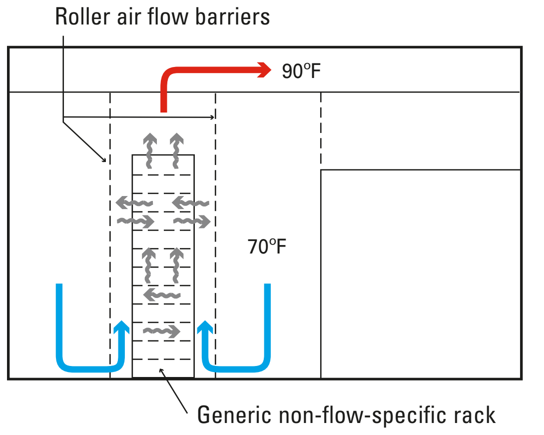

being the ducted aisles), the working (human) spaces would be hot—at temperatures up to ~90˚F—so that error should be avoided.

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.

Issues associated with stratification and pressure balancing inside and across server racks are possible downsides to these approaches. Large or numerous fans are typically needed to maintain the required airflow through all racks. Since racks contain different kinds of equipment and, more notably, different wiring and connections, paths of least resistance to airflows must generally be overcome with fanpower and additional heat/energy loads. However, it may be possible to reduce and help equalize airflow resistance by rearranging equipment or cabling.

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 VFDs[2] 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

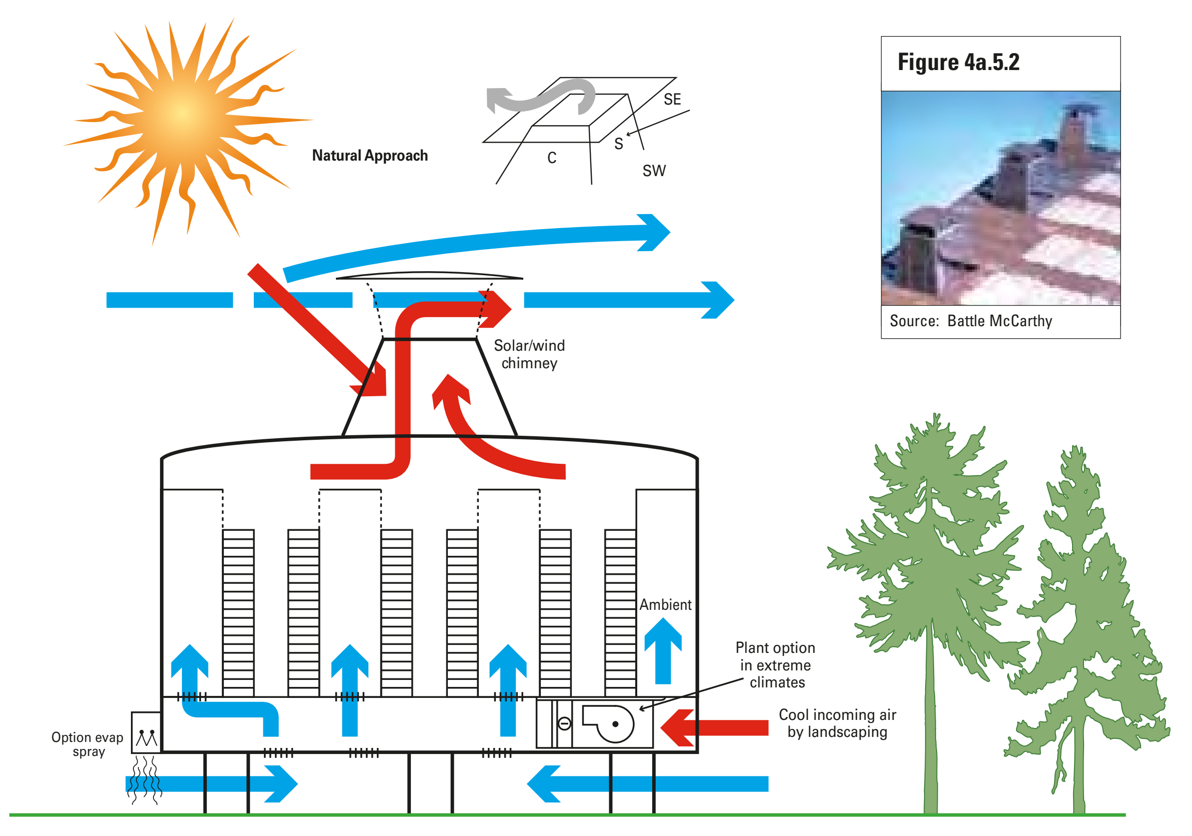

Figure 4a.5.2

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 ACH[6]] 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.

Recommendation

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 (≤200y) in miniature computer fans.
- Keep motors out of airstreams, and buy the most efficient motor on the market (MotorMaster 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.

Recommendation

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

---

[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.

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)*;
- 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.

---

[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).”

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).

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).

---

[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.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.

Recommendation

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.

Recommendation

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.

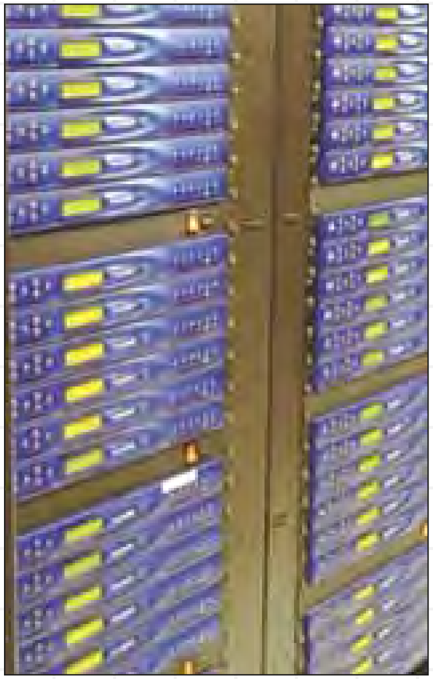

Photo courtesy Chris Hipp

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 65ºF—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 DX[12] 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 65ºF as the practical low setpoint.

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

---

[12] Direct expansion—packaged air-cooled air conditioning units.

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 – 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 unrebutted.

No Straightforward Answer

The typical electric load of a data center (equipment rooms only) is of the order of 400 W/m^2 for internet applications and 200 W/m^2 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

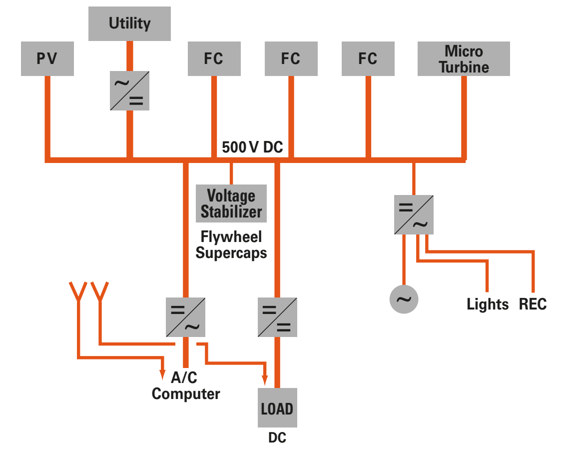

---

[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

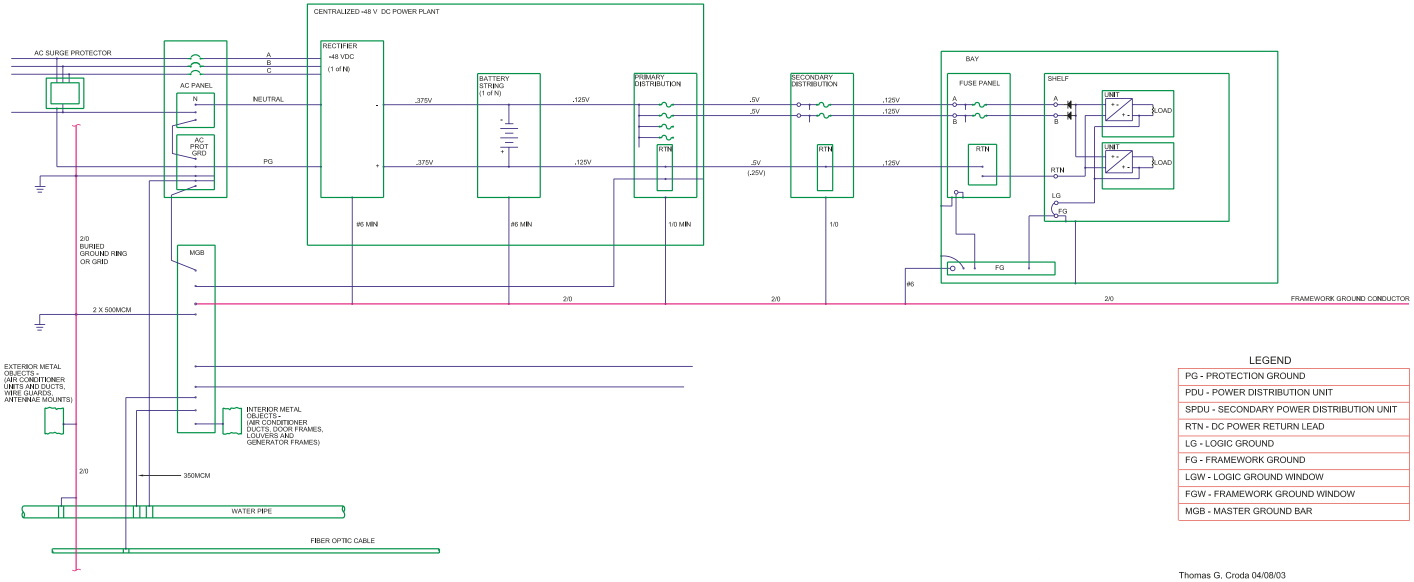

Figure 5.1.3: 540-VDC system

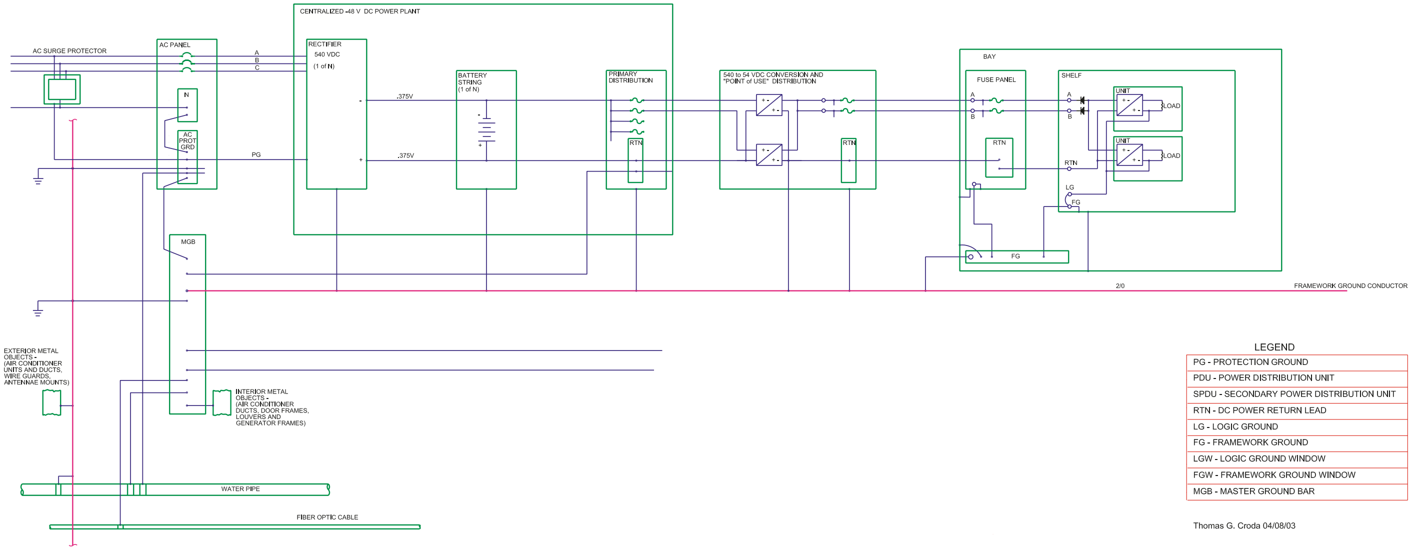

A Preference for DC Power Supply

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

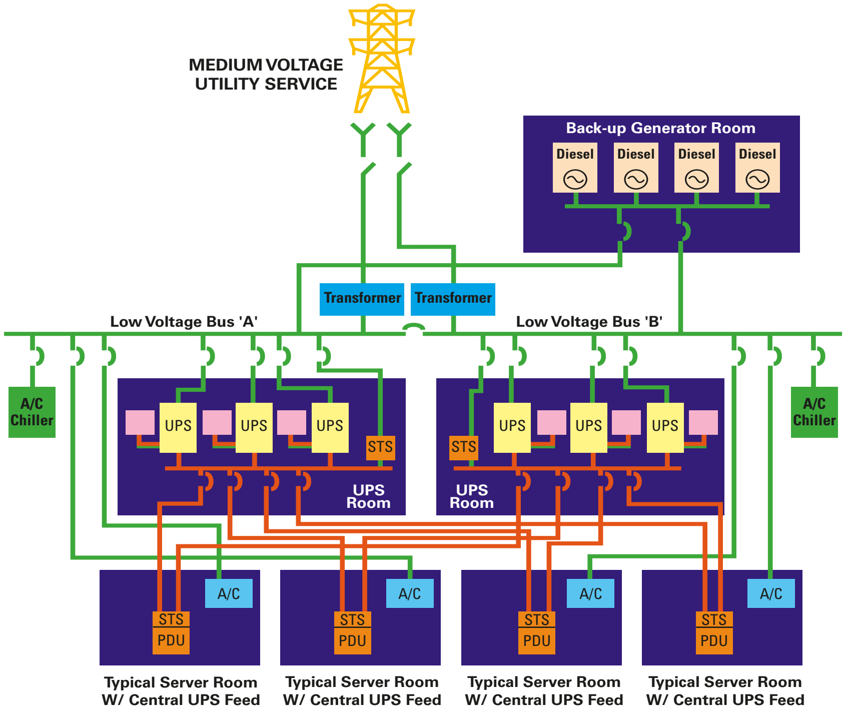

This synergy between the data center's requirement for reliable onsite power and its tremendous—but readily reduced—cooling requirement offers a key strategy for reducing overall power consumption.

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-stabilizing 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.

Figure 5.2.2: Proposed AC system

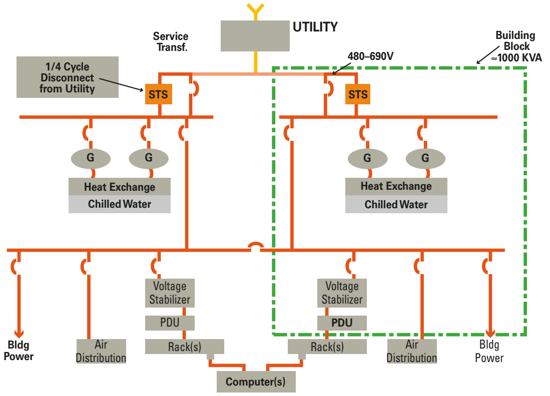

Due to the simpler path, the number and capacity of batteries are significantly reduced, as is the need for breakers and switches. The design eliminates 50 percent of the losses that occur in today's systems.

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.

Recommendation

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 bus[8] 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.

Recommendation

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.)

---

[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).

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.

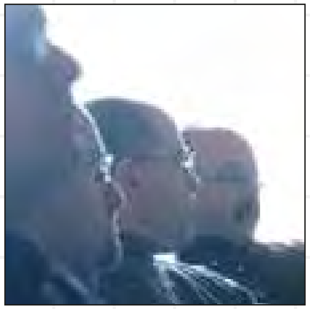

---

[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

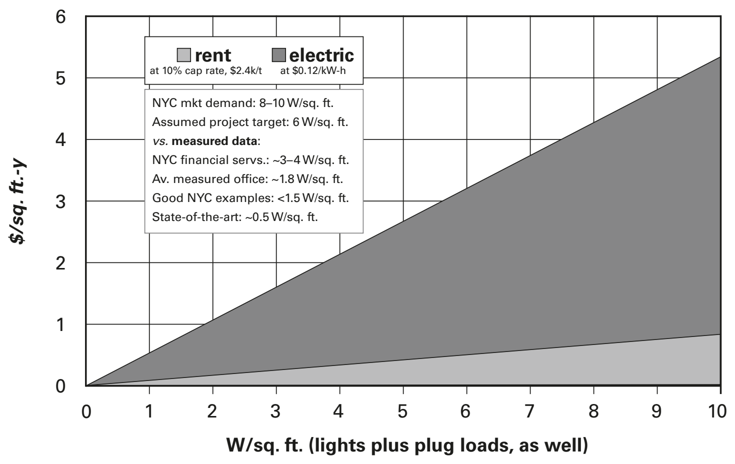

Ken Brill, Executive Director of The Uptime Institute comments: “What is the ‘true’ cost of providing power, cooling, and space for a cabinet of computer equipment? There are two different ways of charging for occupancy and construction. The current method considers only investment cost and operating 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.”

Figure 6b.1: Availability *vs.* reliability
Users should consider tradeoffs of cost, availability, and reliability

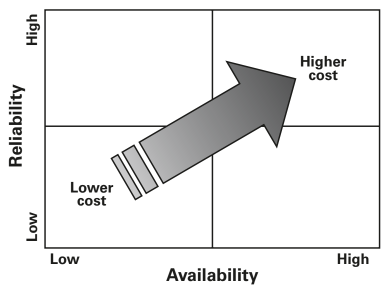

Figure 6b.2: Perception *vs.* reality
Perception and reality are very different with regard to costs

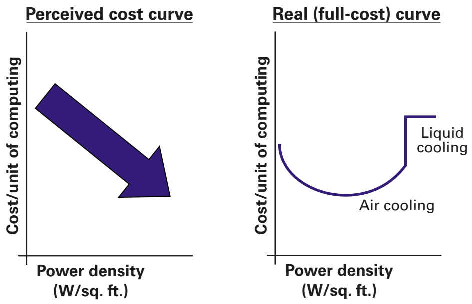

Figure 6b.3: Demand *vs.* price signals
Demand for computing will change with price signals

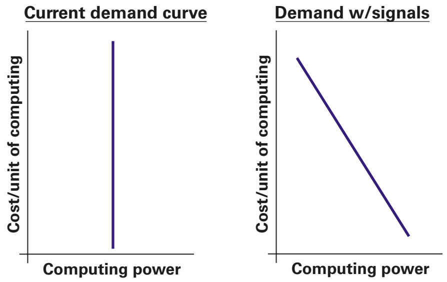

There are myriad disconnects between the narrow foci and missions of the individual sector specialists—real estate, facilities, finance, vendors, IT, and end users—and the best interests of the data center as a whole. As Mr. Brill notes, “Real estate professionals get paid to keep costs low; IT guys 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 analagous 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 output[2]—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/ft^2 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, ft^2/ton, unit of cooling per unit of data-processing;
- Air recirculation—cfm/ft^2, W/cfm, air changes per hour in computer room;
- Power transformer efficiency—percent efficient;
- Lighting—W/ft^2 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."

Recommendation

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 applications[4] 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.

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 executives 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.

---

[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?

Recommendation

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.

Recommendation

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.

Recommendation

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.

Recommendation

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.

Recommendation

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.

Recommendation

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.

Recommendation

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

LEED
LEADERSHIP IN ENERGY & ENVIRONMENTAL DESIGN

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.

Recommendation

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 extended 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.

- 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

UPTIME
TESTED

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 starts 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]

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.

---

[1] www.energydesignresources.com.

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.

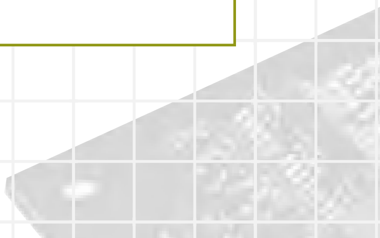

Rocky Mountain Institute (www.rmi.org), founded by Amory and Hunter Lovins in 1982, is an independent, entrepreneurial, nonpartisan, nonprofit applied research center.

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.

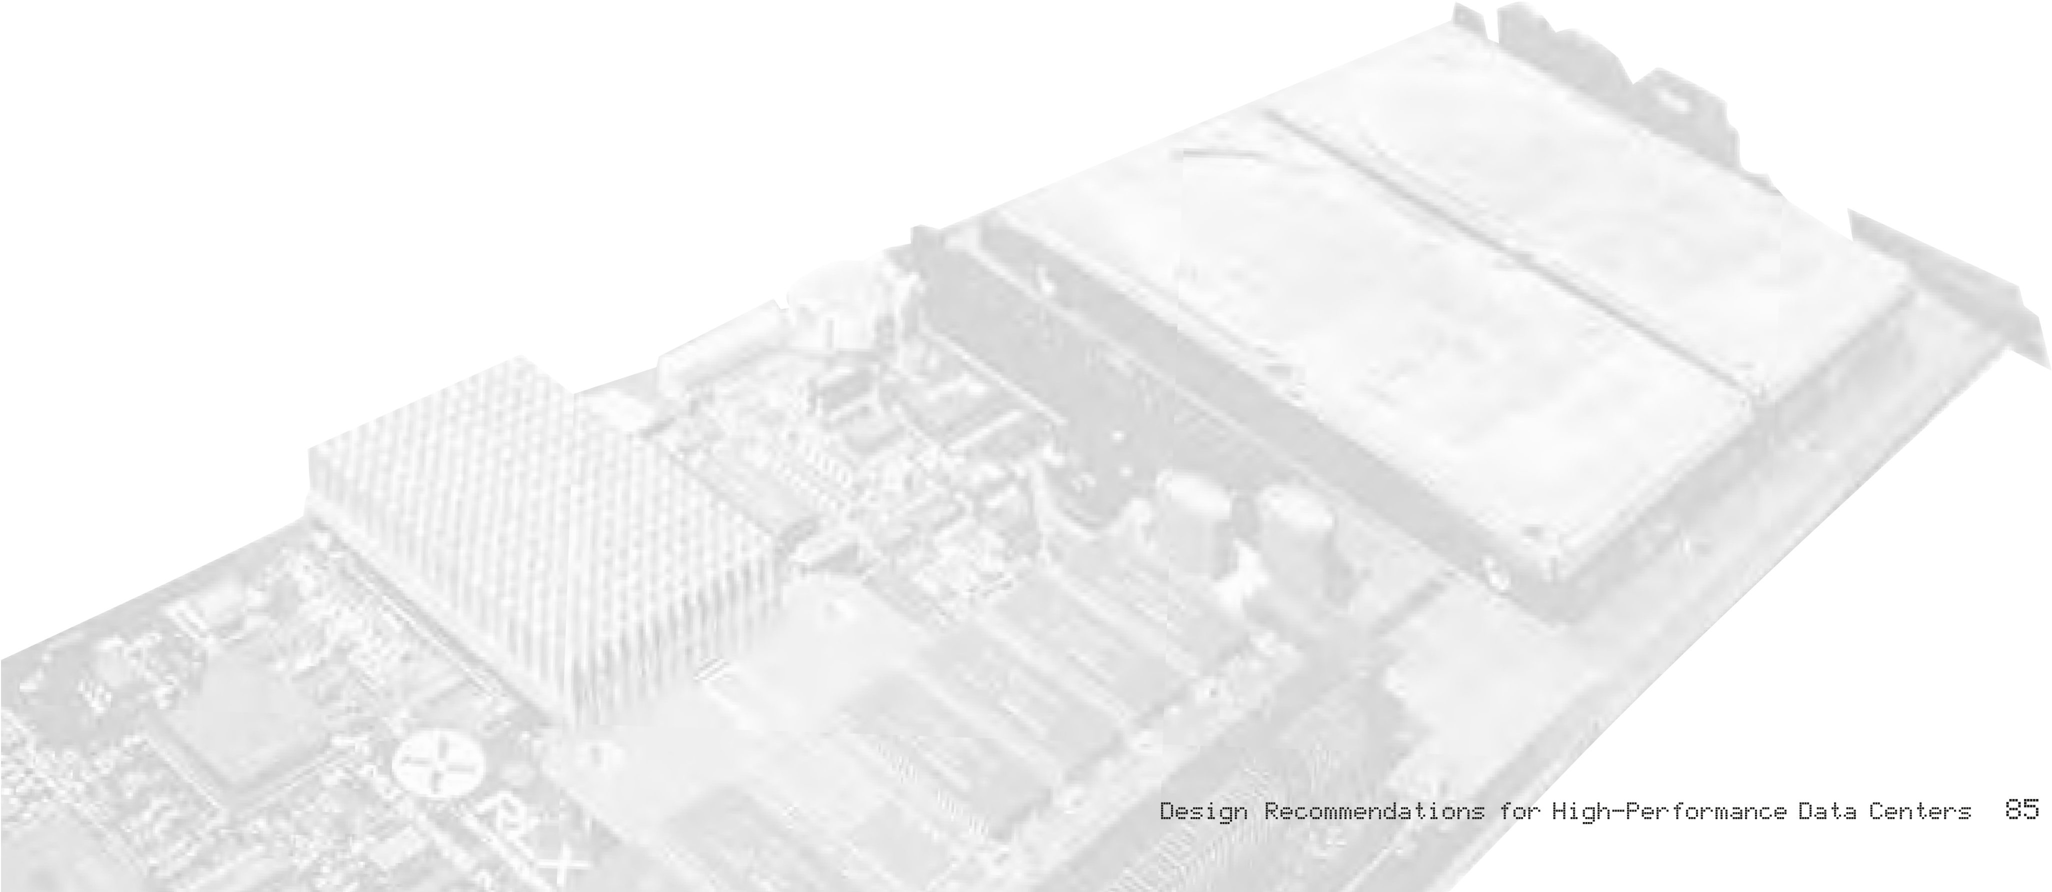

RMI Staff and Consultant team

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/team/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-county 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.cap-e.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.

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+Walhert AG
Leutschenbachstrasse 45, 8050
Zürich Switzerland
+41 305 91 11
www.amstein-walthert.ch
adrian.altenburger@amstein-walthert.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@hippster.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, 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 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
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, 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 Westbrock
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_westbrock@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

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 manage-ment, 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 Cyras (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 flywheel 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.joneslanglasalle.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/cleanroom/) 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

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

General

Pacific Gas & Electric
www.pge.com

California Energy Commission
www.energy.ca.gov

New York State Energy Research and Development Authority (NYSERDA)
www.nyserda.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