REPP

RESEARCH REPORT

Renewable Energy Policy Project

November 2001 ■ No. 14

ENERGY SMART DATA CENTERS: APPLYING ENERGY EFFICIENT DESIGN AND TECHNOLOGY TO THE DIGITAL INFORMATION SECTOR

*by Fred Beck**

* Research Manager, Renewable Energy Policy Project. The author would like to thank Don Beaty, Roger Duncan, Gretchen Parker, Karl Rábago, Keith Reid, Stephen Spinazzola, Dag Tessema, and Chris Wilkins for their valuable input to this work, and Jonathan Koomey, Skip Laitner, Keith Reid, Joe Romm, Virinder Singh, George Sterzinger, and William Tschudi for their timely and insightful review comments. Research for this paper was supported by a grant from the U.S. Environmental Protection Agency. The content of this paper is the sole responsibility of the author, and does not necessarily reflect the opinions of the funding organization, REPP, the REPP Board of Directors, or the reviewers. Fred Beck can be contacted at (202) 293-2898 or via e-mail at fbeck@repp.org.

BOARD OF DIRECTORS

Karl Rábago, Chair
Rocky Mountain Institute

R. Bosworth Dewey
Overseas Trade and Investment Services, LLC

Jeffrey Genzer
Duncan, Weinberg, Genzer, & Pembroke, P.C.

The Honorable Renz D. Jennings
Former Arizona Corporation Commissioner

Alan Miller
Global Environment Facility

David Nemtzow
Alliance to Save Energy

Alan Nogee
Union of Concerned Scientists

Roby Roberts
PacifiCorp

The Honorable Claudine Schneider
Former US Congresswoman

J. Rachel Shimshak
Renewable Northwest Project

Scott Sklar
The Stella Group, Ltd.

Randall Swisher
American Wind Energy Association

Michael Totten
Conservation International

Carl Weinberg
Weinberg Associates

1612 K STREET, NW, SUITE 202
WASHINGTON, DC 20006
202-293-2898
202-293-5857 FAX

http://www.repp.org

CONTENTS

INTRODUCTION **1**
PART I: UNDERSTANDING DATA CENTERS **1**
DATA CENTERS AND THE U.S. ECONOMY **1**
Economic Productivity 1
Economic Importance of Reliable Power 2
ANATOMY OF A DATA CENTER **2**
Typical Layout 2
Box 1: Types of Data Centers *3*
Number and Location 3
Power Quality and Reliability 3
DATA CENTER ENERGY USE **5**
Importance of Choosing a Consistent Metric 5
Box 2: How much energy does the Internet use? *5*
Initial Power Requests Were High 6
Box 3: Data Center Power Requests –Boom and Bust *7*
Measured Power Requirements 8
Five Reasons Actual Requirements are Lower Than Initial Power Requests 8
IMPLICATIONS FOR UTILITIES **9**
Distribution Network Upgrades 9
Lead Times 9
Long-Term Demand Stability 10
Stranded Investments 10
PART II: DESIGNING MORE EFFICIENT DATA CENTERS **10**
ENERGY EFFICIENT TECHNOLOGY AND DESIGN **10**
Server Component Improvements 11
Chiller and Air Conditioning Improvements 11
Air Handling System Improvements 13
Balance-of-System Component Improvements 14
Future Prospects for Improvements 14
EFFICIENCY TARGETS **14**
Efficiency Targets for a Model Data Center 14
National Efficiency Targets for Model Data Centers 16
NEXT STEPS **17**
Box 4: Onsite Generation and Power Park Concepts *18*
Distributed Generation and Combined Heat, Power, and Cooling 18
Renewable Distributed Generation with Photovoltaics 18
Smart Power Parks 18
Rewarding Efficiency 19
Energy Service Companies 19
Industry Cooperation 19
Industry Design Standards 20
Rating, Certification, and Labeling Programs 20
Load Characterization 21
Balance-of-System Component Improvements 21
Zoning Codes 21
CONCLUSION **22**
APPENDIX A: FOOTNOTES AND ASSUMPTIONS FOR TABLES 6 AND 7 ..... **23**
ENDNOTES **24**

ENERGY SMART DATA CENTERS: APPLYING ENERGY EFFICIENT DESIGN AND TECHNOLOGY TO THE DIGITAL INFORMATION SECTOR

By Fred Beck

INTRODUCTION

Data centers, buildings with high concentrations of computers and digital electronic equipment dedicated to hosting websites, supporting e-commerce and providing essential services for the new "digital economy," are a fairly recent phenomena. [1] It was initially reported that power requests of these centers were quite high--one to two hundred watts per square foot or more. This is twenty times greater than modern commercial buildings, which require only five to ten watts per square foot (W/ft2). [2] However, these are most certainly overestimates by substantial margins. More recent data from actual energy demand measurements show data centers draw considerably less power than originally thought.

With recent concerns about electricity supply in parts of the U.S., both utilities and data center owners face challenges in meeting data center electricity requirements with required levels of reliability. [3] While construction of data centers proceeded at breakneck pace between 1995 and 2000 to feed the needs of the expanding digital economy, [4] the bursting of the high-tech stock bubble in 2000 and the 2001 U.S. economic downturn has slowed expansion of data centers considerably. This provides time and an opportunity to examine data center construction and operational practices with an eye toward reducing their energy demands through use of energy efficient technologies and energy smart design practices. As the economy recovers and data center construction resumes, best practices can reduce data center energy use while maintaining or even increasing data center reliability. We find that data center power demands could be reduced by 20 percent with minimal efficiency efforts, and by 50 percent with more aggressive efficiency measures.

PART I: UNDERSTANDING DATA CENTERS

DATA CENTERS AND THE U.S. ECONOMY

Data centers, while comprising a very small fraction of the U.S. building stock, are becoming increasingly important for the U.S. economy. They are an essential component of the infrastructure supporting the Internet and the digital commerce and electronic communications sector. Continued growth of these sectors require a reliable infrastructure because, as described below, interruptions in digital services can have significant economic consequences.

Economic Productivity

The digital economy, based on information and communication technology (ICT), is rapidly becoming a significant part of the U.S. economy. From 1990 through 1997 the nation's economy as a whole averaged 2.6 percent growth per year, while the ICT sectors—excluding electronic sales—averaged 13.5 percent growth over the same period. [5] While ICT industries still account for a relatively small share of the economy's total output—an estimated 8.3 percent in 2000—they contributed nearly a third of real U.S. economic growth between 1995 and 1999. [6] For example, in 1999 the internet supported $300 billion in electronic commerce revenues. This is predicted to grow to $1.3 trillion in business to business commerce revenues by 2003. [7]

As recently as September, 2000, it was estimated that worldwide spending for ICT products and services will reach $2.6 trillion in 2005. [8] Even with the recent economic slowdown, continued growth in this sector is expected. Electric power reliability and availability for data centers that host internet

sites and manage digital data transfer, storage, and calculations will only become more important.

Economic Importance of Reliable Power

Power outages and power quality issues can cost U.S. businesses dearly. [9] Power failures are particularly vexing for business that rely on digital circuitry, for a power interruption or power quality drop of just a fraction of a second is enough to crash computers or halt industrial processes relying on digital process control circuitry. For example, credit card processing centers can lose over $2.5 million per hour from power interruptions, and computer chip fabrication plants can lose up to $60 million per power interruption. [10] The importance of reliable power for the industrial and digital sectors is further underscored by a report recently released by EPRI's Consortium for Electric Infrastructure for a Digital Society (CEIDS). [11] EPRI reports that U.S. digital economy firms lose $13.5 billion to electric outages annually, primarily from lost productivity and idled labor. [12] Inclusion of continuous process manufacturing, fabrication, and essential services such as transit, water, and gas industries increases the annual U.S. economic loss to $45.7 billion from power outages plus another $6.7 billion from power quality issues. Finally, when all business sectors are included, the U.S. economy is losing between $119 billion and $188 billion annually due to power outages and power quality issues.

ANATOMY OF A DATA CENTER

Data centers house a high density of digital electronics and computer technology, requiring higher quality and more reliable electric power than most commercial buildings. They are essentially building shells packed with computers, power supplies, power conditioning equipment, control electronics, and backup power systems along with air conditioning systems to keep the equipment cooled to optimum operating temperatures, generally 68-70 °F. The computers used in data centers are generally known as servers. Multiple servers are secured in racks that typically have a 2 foot by 2 ½ foot footprint and are 70 to 87 inches high. These racks are placed on a raised floor area, which serves as a plenum allowing cooled air to move below the racks, then up through perforated floor tiles to cool the racks before being drawn back through the heating, ventilation, and air conditioning (HVAC) system. Power supplies, conditioning equipment, and backup generators are placed separate from the raised floor area.

Data centers range in size from a small computer room housing a few server racks to 200,000 square feet (ft2) or greater dedicated facilities holding tens to hundreds of server racks. One of the largest proposed data centers is a 2.2 million gross ft2 U.S. DataPort data center complex, which recently received permitting for construction in San Jose, California. [13] Because data centers can have different types of tenants and ownership structures, a number of different terms are used when referring to these facilities. Some of the common terms are given in Box 1 (following page).

Very few people work in a typical data center, unless the building housing it also provides office or retail space. Due to the proprietary nature of digital data and tight competition within information technology markets, data centers are designed for maximum security. Many have no exterior windows, some employ bullet-proof glass partitions and Kevlar reinforced walls. They can be purpose-built as new construction, or retrofit from existing buildings. Retrofit buildings include office buildings and warehouses, as well as industrial facilities which have high ceilings and greater load-bearing design, which facilitates the routing of cables and installation of heavy HVAC and power generation equipment.

Typical Layout

Because data centers vary greatly in size, function, and building configuration, it is difficult to define the layout of an "average" data center with any precision. However, a number of publications have provided general guidelines of space utilization in data centers. Table 1 describes area usage of a generic data center as a whole, and Table 2 (page 4) provides a breakdown of the raised floor area usage. In this example, the actual computer hardware (server racks) accounts for only 12 to 15 percent of the total building area. (I.e., half the building is

Box 1: Types of Data Centers

"Data center" is often a generic term used to describe a number of different types of facilities that house digital electronic equipment for internet site hosting, electronic data storage and transfer, credit card and financial transaction processing, telecommunications, and other activities that support the growing electronic information-based economy. There are many flavors of data centers--each may have its own specific issues related to power demand and load profiles. A few common terms and client profiles for data centers are:

- **Data storage and Internet hosting facilities**, also known as server farms or internet hotels, that perform a variety of functions.
- **Internet Service Provider routers** (ISPs), dedicated specifically to supporting the Internet.
- **Telecommunication switches**, known as telecoms or telcos. These are more energy-demanding than typical Internet data centers. [14]
- **Corporate data centers**, where racks and computer equipment are wholly owned and operated by the corporation. Because one entity owns and operates this type of data center, energy demand may be better-characterized than other data centers. Other examples of this type of data centers may be found in universities and other institutions.
- **Managed data centers**, where racks and computer equipment are owned by data center owner, but leased to tenants.
- **Co-located server hosting facilities**, also known as CoLos, where rack space is leased by tenants and computer equipment is owned and operated by tenants. Because tenants may move in and out, upgrade their computers frequently, and have a disconnect between the energy-using facility and the billing department, energy demands tend to have greater fluctuations and to be less-well characterized than corporate data centers.

raised floor area, but only 25-30% of this floor area is actually occupied by the computers. Thirty percent of fifty percent is fifteen percent.)

Table 1 Approximate usage of total building floor space for a dedicated data center

Area	Usage (%)
Core and common area (halls, stairways, elevators, electrical and mechanical space)	20-25%
Generators, batteries, power supply and conditioning equipment	25-30%
Raised floor area	50-60%

Source: Thomas Callsen. "Internet Hotels--A Big Draw" Transmission & Distribution World. February 2001. And personal communication with Keith Reid, PG&E, November 8, 2001.

Number and Location

Data centers are generally concentrated in cities with major fiber optic nodes (Figure 1, following page). For example, at the end of 2000 there were approximately 320 Internet Service Provider (ISP) data centers in the US alone with total computer room space occupying 9.5 million square feet - an average of some 30,000 sq.ft. per data center. An August, 2000, report from Salomon Smith Barney anticipated that there would be 17.9 million square feet of data center rooms built or under construction in the U.S. by the end of 2001, an 80 per cent increase over 2000.[15] However, current economic conditions have undoubtedly slowed or even temporarily stopped this expansion. [16]

Power Quality and Reliability

The marketability and commercial success of some types of data centers rests on their ability to assure the customer that they have the highest possible power quality and reliability, 24 hours a day, 365 days a year. Data centers go to great lengths to

Table 2: Approximate usage of raised floor space area

Area	Usage (%)
Computer hardware footprint (servers, rack-mounted equipment, telecommunications frames)	25-30%
Service clearances around products (allows for movements of cooling air and personnel)	30-35%
Infrastructure support equipment (in-room uninterruptible power supplies (UPSs), power distribution units (PDUs), cooling systems, air handling equipment, and other support electronics)	20%
Main aisles, support columns, other non-electrified areas	20%

Source: Uptime Institute. "Heat Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment" 2000.

ensure reliable supply, often requesting that electric utilities run two or more separate feed lines to data centers to improve power reliability and lower the chance of power interruption. Power quality is ensured with extensive power conditioning equipment, with back-up power systems including batteries, flywheel storage, ultra-capacitors, and on-site power such as diesel generators ensuring supply in case of a power failure. In addition, many data centers have multiple redundancies built into their power and HVAC systems to ensure that power and thermal management are maintained at all times. Even a brief power outage can lead to significant revenue loss, as well as very costly equipment damage if air conditioning and cooling systems are not brought back on line quickly.

While power reliability is often discussed in terms of percentage availability throughout a year, the number of times power is lost may actually be of greater relevance to data centers. Power supplied by utilities is generally available 99.9 to 99.99 percent of the time. This is referred to as "three-nines" to "four-nines" of availability. Much of the data center literature suggests 99.9999 percent availability, or "six-nines", is required by data centers, hence the need for backup batteries and on-site backup generation. [17] However, on-site human actions are often more likely to cause power interruptions than are utility interruptions. The Uptime Institute states that due to fire and electrical safety codes, fire alarm tests or Emergency Power Off (EPO) procedures can contribute to power interruptions. Assuming only

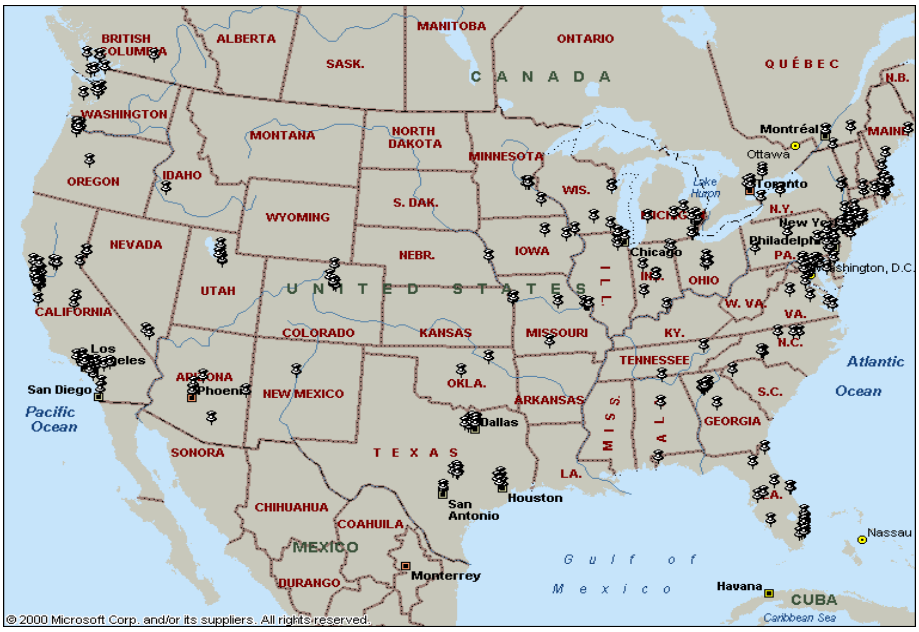

Figure 1. Map of U.S. Colocation Facilities
Reprinted with permission from www.colosource.com

one fire alarm or EPO in a five-year period, Uptime finds that the best possible reliability for data centers is 99.995%, or about 26 minutes of downtime per year. [18] So, while data centers may pursue a goal of six or more nines of power reliability, attention must be paid to factors other than the reliability of the utility connection and data center hardware.

DATA CENTER ENERGY USE

Data center power requests based on measured data are much lower than requests based on design specifications. Data center owners previously specified power requirements of 100 to 200 watts per square foot of raised floor area (W/ft2) and greater. Rapid construction of data centers over the past five years or so to meet quickly increasing and highly competitive demands for digital information services has led to inefficiencies in data center energy use. However, data center owners and operators do not want to release their power requirements or load profiles publicly because they fear this could hurt their competitiveness. Lack of clear definitions of power demand make data center energy characterizations difficult, while a number of factors lead to overstated power requirements. The power requirements of the Internet as a whole, while now better understood, were also initially overstated (see Box 2).

This section explains why choosing a consistent definition of power density is important for understanding data center energy use, gives examples of energy requests before the dot.com bust and 2001 economic downturn, provides best estimates of actual energy requirements, and explains five reasons data center energy requests are larger than needed.

Box 2: How much energy does the Internet use?

Initial estimates of Internet energy use were overstated, and have since been shown to be inaccurate. However, because these high numbers continue to find their way into print and be used by some data center industry analysts, we take a moment to dispel any remaining confusion.

The source of initial high estimates was a May 1999 article by Mark Mills and Peter Huber published in Forbes magazine. The article stated that electricity use associated with the Internet totaled about 8 percent of all U.S. electricity use in 1998, and would grow to require half of all U.S. electricity output in the next decade.[19] Subsequently, the U.S. Environmental Protection Agency requested that scientists at Lawrence Berkeley National Laboratory (LBNL) examine the numbers in the article to determine if they were correct, and if not to provide a more accurate estimate. The LBNL study found the Forbes article overestimated Internet electricity use by a significant margin. [20]

As it turns out, Mills and Huber's analysis contained several errors, including energy use calculations for major dot-com companies, Web servers, Internet routers, telephone switches and home and business PC's. [21] For example, Mills and Huber assumed mainframe computers require 250 kW apiece for operation and cooling, while LBNL measured energy requirements of 19.2 kW apiece for mainframe operation and cooling--thirteen times less energy than was assumed. [22] In the end, LBNL found electricity use associated with the Internet in 1998 was 36 million MWh, eight times lower than Mills and Huber's original estimate of 295 million MWh. [23]

Importance of Choosing a Consistent Metric

While there is a great deal written about data center energy use, comparisons of energy requirements between them are often difficult or confusing because authors do not always define their energy requirements in a consistent manner. This creates headaches and uncertainty for utilities, who need accurate information in order to supply the correct amount of electricity to the data center customer, as well as to gauge whether or not new generation or distribution capital expenses need be incurred to meet these requirements. Utilities and data center operators often quantify power requirements in terms of power density, or watts per square foot of floor area (W/ft2) that a building will use. Problems arise when different parties use different definitions of floor area (some may include office spaces, others may not), and different sets of energy-using

components (some may include energy-using equipment outside the computer room, others may not) to define this power density. These issues have been discussed in detail in Jennifer Mitchell-Jackson's May 2001 thesis, "Energy Needs in an Internet Economy: A Closer Look at Data Centers." Building on this work, at least four different metrics of power density can be defined (see Table 3). [24]

Table 3. Differing definitions of power density

Rack (or product) power density	Power drawn by an individual rack (or product such as a server or tape backup device), divided by it's footprint
Simple computer room power density	Power drawn by the computer equipment installed on the raised floor, divided by the raised floor area that houses the equipment
Total computer room power density	Power drawn by the computer equipment and all supporting equipment such as power distribution units (PDUs), uninterruptible power supplies (UPSs), HVAC, and lighting, divided by the raised floor area
Building power density	Total power drawn by the entire building, divided by the total floor area of the building

- **Rack power density** is generally the highest because a rack packed full of energy-intensive computer equipment draws the greatest load per square foot of floor space of any component within a data center.
- **Simple computer room power density** is lower than rack power density because the energy required by the racks is spread out over the entire raised floor space, including aisles and non-energized areas as defined previously in Table 2.
- **Total computer room power density** is higher than simple computer power density because it includes the power requirements of all equipment that supports the operation of the data center.
- **Building power density** generally the lowest because it includes floor space that is much less energy-intensive than the computer and support electronics rooms themselves, such as hallways, closets, and standard office spaces.

Total computer room power density is most representative of the true power requirements of data centers because it accounts for all direct and indirect computer energy use in a data center, assigning it to the raised floor area where the computer racks are housed. It can be thought of as the "fully loaded" power requirement per square foot of computer room space. Total computer room power density is also the most useful for comparing data centers, as it allows comparisons between buildings of different sizes and office space usage. For mixed-use buildings, utilities will often define a separate power density for the commercial office space or warehouse areas of the building for use in making load projections. [25]

Initial Power Requests Were High

Data center operators requested high power densities to satisfy their customers that they have the power capacity on hand to be competitive, though in doing so they request much more than they actually require. [26] Extensive discussions with data center industry representatives reveal that the overriding motivational factors for data center operators are profit (often expressed in $/ft2) and time to market. When questioned about energy use, almost all asked said the cost of energy to run the data centers is so small compared to profits that it is not even on their radar screen, and that virtually no data center operator even stops to consider employing energy efficient design or technology. [27] However, this may be less true for corporate data centers that have a longer time horizon than many Internet-based data centers, [28] and power costs may become more relevant due to recent economic slowdowns as well as potential electricity price increases. Data center customers demand that power capacities are sufficient to supply a fully built-out data center (all racks filled to maximum server capacity), with the assumption that power

draw will equal the nameplate rating of each component (maximum possible power draw).

Many sources put data center requests to utilities at total computer room power densities between 100 W/ft2 and 200 W/ft2, and in cases up to 300 W/ft2. [36] A March, 2001, survey by a major IT website indicated that 85% of data center operators polled maintain power capacity to feed a total computer room power density of at least 100 W/ft2, and over 20% of all operators polled maintain capacity to feed power densities of 200 W/ft2 or more. [37] By comparison, commercial office buildings typically draw 5-10 W/ft2, or ten to forty times less than power densities being requested by typical data centers. [38] In calculating total building load, we must remember that the power density (in W/ft2) of data centers is based on the raised floor area. For example, if we assume a 200,000 ft2 data center has a total computer room power density rating of 100 W/ft2, and 50% of the data center building is devoted to the raised floor area, the data center load will be 10 MW. (I.e., 100,000 ft2 * 100 W/ft2 = 10 MW.) This is enough power to supply approximately 7,500 average U.S. homes. [39]

Before the dot.com crash and 2001 economic downturn, high requested power densities coupled with the expansion of the information and communications technology sector led to a number of very large requests for power from utilities. These requests as reported in the media grabbed utilities attention with regard to potential data center energy use. Box 3 gives examples of these requests.

Box 3: Data Center Power Requests—Boom and Bust

Prior to the 2000-2001 dot.com crash and 2001 economic slowdown, a burgeoning data center construction industry placed increasing power demands on local utilities, often with much shorter lead times on power delivery requests than utilities were used to for industrial-sized loads. Below are examples of data center electricity requests as then reported by utilities and the media:

- In December 1999, Seattle's Puget Sound Energy (PSE) had no requests for electricity to serve large Internet data centers. In August 2000, they had requests for 445 MW from such centers in a small area near Southcenter Mall. PSE's Karl Karzmar stated that this power draw "…is the equivalent of six oil refineries." By September 2000, requests had reached 700 MW, briefly rising to over 1,000 MW until the dot.com bust came. As of July 2001 total requests remained at about 750 MW, which over a 4 year planning horizon would more than double planned PSE loads. [29]
- In June 2001 the New York Times reported that just one new server farm, proposed for the economic development zone of the South Bronx, would draw more than twice as much power as the entire former World Trade Center complex. In April 2001 Consolidated Edison reportedly warned that plans for 46 such data centers proposed for New York and Westchester County over the next four years could impose another 500 megawatts of demand (an increase of 4% over current load) on its already strained system. [30]
- The New York Times also reported in July, 2000 that a developer looking at a site in North Jersey for the home of a potential million square foot data and communications center requested 100 MW of power from Public Service -- one third of what the utility provides to the whole city of Newark." [31]
- According to a November 2000 Sacramento Bee article, one data center company told the Sacramento Municipal Utility district that it would need 50 to 65 MW of power--roughly the equivalent of all other growth in the area in an average year.[32]
- Keith Reed, a Senior Corporate Account Manager with PG&E, said that data center customers in PG&E's territory are forecasting unheard of leaps in electrical demand. In October 2000, PG&E reported that data centers requested 341 MW of power in 2000, and an additional 1000 MW of power by 2003--the equivalent of approximately three new power plants.[33] As it turns out, only 50 MW of demand has been realized as of November 2001. [34]
- In Texas Austin Energy had planned to provided 100 new MW to data centers in 2001. However, as of November 2001 only 6 MW of new load has materialized. [35]
- As illustrated by the cases of PG&E and Austin Energy, data center power requests have dropped dramatically in recent months, giving utilities time to work with data center developers to ensure that a load crisis (or apparent crisis, at least) does not happen again.

Measured Power Requirements

Though detailed, validated energy use data is not readily available, a literature search and interviews with industry representatives indicate that in practice data centers draw approximately 20-50 W/ft2, [40] which is generally one-fifth to one-third of data center capacity requests in 2000. A summary of these load measurements and estimates is given by Table 4. Because actual loads are significantly lower than design specifications, energy inefficiencies are likely to result from over-sized HVAC and power electronics. In addition, customer loads that are significantly different than originally projected can upset utility load planning. These issues are discussed further in following sections.

Five Reasons Actual Requirements are Lower Than Initial Power Requests

Five basic assumptions about data center energy intensity are made in planning power requirements for data centers that lead to an initial overestimation of the power demand of server, HVAC, and balance of system equipment. [48] This in turn leads to designing oversized HVAC and balance of systems components, including power control and distribution electronics, as well as back-up on-site generation capacity. Standard engineering practice then adds design safety margins on top of this, leading to the result seen above, actual loads a third or less of original design loads.

Nameplate Power Assumption

The first assumption made is that all server components draw the amount of power specified on the nameplate of that component. The nameplate rating is the maximum possible power demand of the component, usually measured in amperes. For safety reasons, most computer equipment never draws more than 80% of the rated power even at peak demand. [49] In fact, research has shown that most PC's draw only 20-35% of their nameplate rating. [50] For example, a Sun Ultra server rated at 475W used 142W at startup and 133W under normal use, only 30% and 25% of its nameplate rating, respectively. [51] Thus, basing power requirements on nameplate ratings significantly overstates power requirements.

Fully Loaded Server Assumption

The second assumption is that each server is fully

Table 4. Measured Data Center Energy Demand Data from Various Sources

Source	Total Computer Room Power Density (W/ft2)	Comments
Infomart [41]	50	Based on six months of monitoring by a Texas-based data center developer
Energy Information Administration [42]	50	Annual Energy Outlook 2001
Jennifer Mitchell-Jackson and Lawrence Berkeley National Laboratory [43]	35-50	Measured energy demand of U.C. Berkeley data centers
Commonwealth Edison [44]	40	Utility data
Pacific Gas & Electric [45]	35	Utility data, fully occupied facility, actual load 33% of requested capacity
Jennifer Mitchell-Jackson: Case Study [46]	32	Detailed case study, actual load 25% of design load
Edison Electric [47]	25-40	Utility data, requested capacities 150-200 W/ft2

configured. This means that every expansion card slot is filled, the maximum number of processors supported is installed, the maximum amount of memory is installed, the maximum number of hard drives are installed, and every input and output port is in use. It is rare in practice to find servers configured this way, most are only partially loaded, and therefore draw less power than a fully loaded server.

Fully Loaded Rack Assumption

The third assumption is that each server rack in the data center is fully filled with servers. Typical modern server racks can hold up to 42 "1U" servers. (A "server unit", or "U", is 1.75 inches in height). However, a recent survey showed that the average rack in a sample of data centers was only one-third filled with servers. [52] A partly full server rack draws less power than a fully loaded one.

Fully Built-Out Data Center Assumption

The fourth assumption is that the raised floor area will be fully populated with server racks. This may not be the case. Data center companies often purchase or lease more space than they immediately need so that they will have room for expansion. A recent Salomon Smith Barney report stated that data center revenues are usually calculated based on the data center being filled to 30-40% of capacity. [53]

Balance of System Assumptions

Power supply and conditioning electronics, HVAC, backup generation, and other balance of system components are sized based on the previous four assumptions being true--that server components draw full nameplate load, that servers are fully configured, that all racks are fully filled and in use. In addition to this, engineers will typically add another 10-20% of safety margin capacity. Due to the synergy of building loads, over-sizing non-HVAC balance-of-system components adds an additional heat load, so HVAC must be further oversized to compensate for this. Finally, it will be assumed that all balance of system components will draw full nameplate load, which maximizes the amount of power a data center will request based upon the desire to never risk running out of power to run the entire facility at maximum load.

IMPLICATIONS FOR UTILITIES

Large power requests coupled with much smaller realized power requirements have focused utilities attention on better understanding data center power requirements. The large and relatively constant loads of data centers are very attractive to utilities from a revenue standpoint. However, utilities often have concerns about meeting data center load requirements and time frames, as well as about the long-run stability of the load and the revenue it brings. Additionally, discrepancies between requested and measured power requirements makes utility load planning difficult, and subjects utilities and ratepayers to potential stranded investments.

Distribution Network Upgrades

First, difficulties may be encountered when expanding or retrofitting existing power distribution grids to handle loads from new data centers, particularly in downtown locations or suburban areas. For example, a new data center may request 40 MW on a site originally planned to draw 4 or 5 MW. In order to accommodate this load, distribution companies must upgrade their distribution network and install additional feeders, and in many cases build new substations to serve the load. Substations can be particularly expensive in downtown locations. For example, costs for a substation in urban Chicago have been estimated at $60 million, 5 times that of a rural substation installation. [54] In some cases, system-wide grid upgrades may be necessary to service data center loads. [55]

Lead Times

Second, data centers can be sited quickly and may require power from the utility within two to six months of their first request. Because time-to-market is a very strong factor for the commercial viability of data centers, they are likely to site their facility elsewhere if the utility cannot deliver on time. This is challenging for utilities, as traditional lead times for installation of new generation, transmission or distribution capacity are traditionally on the order of one to three years. If they cannot deliver within the data center's time frame, they miss out on a lucrative opportunity.

Long-Term Demand Stability
Third, data centers, and in particular co-located facilities, differ from other high power density commercial and industrial loads because they tend to be much more transient customers. The competitive nature of the information technology business means that data center companies pop up and disappear on a regular basis. While a data center or its tenants may sign a three-year contract with a utility, it may be gone long before that as a result of short term economic or industry trends. Semi-conductor manufacturers don't just pick up and move to a new location, but a data center can do that very easily, even abandoning obsolete technology at their current site in favor of newer, more competitive technology in a new location, potentially outside the utilities service region. Finally, new digital technologies may increase or decrease load requirements on a fairly quick time-scale.

Stranded Investments
Fourth, utilities are faced with emerging data indicating that data centers typically draw only a third of the load they request. This makes load planning and capital expense recovery difficult. In addition, as shown previously in Box 3, worsening economic conditions can reduce data center power demands drastically. Electric utilities are reluctant to invest significant sums of money on local distribution network upgrades without assurances that the load will develop and will be in existence in five to ten years or longer so that capital costs can be fully recovered.[56]

Utilities deal with these issues in different ways. For example, Austin Energy has robust markets for future electricity growth regardless of data center demand, and are focusing on accurately predicting load profiles of data centers rather than writing legal agreements to deal with potential stranded costs. Other utilities are installing tariff structures or zoning requirements to address data center contingencies. Chicago's electric company, Commonwealth Edison, recently imposed a requirement that data center operators put up large deposits - in some cases as much as $10 million and more - to gain service for a facility. Florida Power & Light has also begun seeking deposits from carrier hotel operators to ensure that they are not stuck with stranded investments.[57]

PART II: DESIGNING MORE EFFICIENT DATA CENTERS

ENERGY EFFICIENT TECHNOLOGY AND DESIGN

Energy-efficient technologies, smart design and proper preventative maintenance procedures can significantly decrease data center energy demands. During the time of rapid data center expansion, operation of such centers was very lucrative, with revenues estimated as high as $1 million dollars per minute. [58] According to many people in the internet-hosting industry, getting a new data center on line as fast as possible (i.e., minimizing time-to-market) was the overriding factor in designing and building data centers. [59] However, haste to install data centers, lack of accepted and standardized design guidelines, and lack of financial incentives to save energy [60] has led to inefficient design, poor design implementation, and use of energy-inefficient technologies in data centers.

When actual building power loads are much less than design loads, energy inefficiencies often occur. For example, standard commercial building energy audits commonly show that oversized air conditioning systems waste energy. A switch to smaller compressors which better match the load, the addition of variable speed controllers for ventilation fans, and microprocessor-based sensors and process controls commonly reduce HVAC energy consumption by 40%, often with very short payback times. [61] Because actual data center loads are significantly less than their design specification, similar opportunities to reduce power use exist.

As in standard commercial building efficiency practice, a data center must be treated as a whole to exploit synergies between design, technology, and operation if maximum efficiency and reliability are to be realized. For example, internal building heat loads from servers, power supplies, and the like

should be reduced before re-sizing the HVAC system for maximum efficiency, and maintenance routines should be enforced that keep the HVAC system running at peak efficiency once it is in place. Data center operators place a high premium on power reliability and quality, and tend to be very concerned that changes in technology or practices from what they are already using will reduce reliability. Thus, energy efficient technologies and improved management practices that increase reliability and/or power quality while saving energy are more likely to be considered than ones that do not offer increased reliability. Further, solutions that are modular and scalable will benefit data centers because they can be expanded as the data center builds out, reducing upfront capital cost, and offer consistency between data centers of different sizes, reducing the need to maintain many dissimilar systems across facilities.

Server Component Improvements

Processor chip power requirements and number of chips within a server effectively determine server power draw. This then determines the amount of waste heat that must be rejected from the server in order to ensure long processor lifetimes and high reliability. Efforts to make smaller and more energy efficient processors have accelerated with the boom of laptop computers and the need to extend battery life and reduce internal heat generation by minimizing power draw. As an example of recent progress, while Intel's recent 1.8 GHz Pentium IV processor consumes up to 66 watts [62] and earlier Pentium III processors required 33 watts, Intel announced in November, 2000 plans to market a new Pentium III processor that requires less than half a watt (0.5 W). [63]

Recognizing the need to reduce computer heat output through improved energy efficiency, IBM recently established a low-power computing research center in Austin, focusing on increasing the power efficiency of computer chips, servers, software and other computing systems and devices. IBM believes new chip design techniques under development can lead to devices that operate on one-tenth the power of current devices. The U.S. Environmental Protection Agency (EPA) recently awarded an ENERGY STAR "Certificate of Recognition" to IBM's eServer z900 mainframe for its contributions to energy savings. It is the first server to receive such an EPA award. [64] In June 2001 EPA also awarded an ENERGY STAR Certificate of Recognition to Transmeta for its energy-efficient Crusoe microprocessor, which uses half a watt (0.5 W). EPA also noted that Transmeta's innovative approach to microprocessor design is enabling many EPA ENERGY STAR Partners to meet program energy-efficiency specifications. [65]

While more efficient processors can save energy, the trend is to pack more and more processors into a server, creating "ultradense" servers, also known as "server blades." Manufacturers need to be careful to avoid increasing overall power requirements and heat output with these ultradense configurations, or they are likely to exceed the cooling capacity limits of conventional data center cooling systems. Fibercycle is now using the Crusoe processor in its WebBunker server rack product, which can house up to 504 processors per rack verses 40 to 84 processors in conventional racks. Because the Crusoe chip is so much more efficient, Fibercycle claims their product requires one-quarter the power of competitive products. [66]

Finally, because not all servers are needed at all times, Tadpole and Platform Computing have introduced their "PowerBack" technology which monitors and powers down unneeded servers, powering them back up when required. They estimate that PowerBack could reduce energy consumption by 15 percent. [67]

Chiller and Air Conditioning Improvements

Roughly 40 % of the power used annually at Internet server farms is for air conditioning and ventilation to keep the racks of computers from overheating and breaking down. [68] In a detailed case study by Jennifer Mitchell-Jackson, HVAC loads accounted for 37% of total computer room power draw, while PG&E comes in with a lower estimate of 25-31% of total building load. [69] As previously mentioned, part of the reason for high HVAC power draw is that these systems are often greatly oversized compared to actual cooling requirements.

Running HVAC systems at part load is very inefficient. Another component of high HVAC power draw is use of conventional rather than highly energy-efficient cooling technologies. It has been estimated that poorly designed HVAC systems in data centers use twice the energy of a well-designed efficient system. [70] One PG&E representative suggests that cooling loads could be reduced to 15-20% of total building load. [71]

There are a number of commercially available off-the-shelf technologies that can be employed to improve data center HVAC efficiency:

- At 0.5-0.7 kW/ton, central chilled water systems are approximately twice as efficient as direct expansion air cooled systems that draw 1.2 kW/ton. [72] Concerns that water near computer equipment may cause problems can be ameliorated through proper chilled water system design.
- Economizer cycles can be used with central water chilled systems or water-cooled direct-expansion chillers to cool chiller water almost for free when wet-bulb temperatures are below 40 °F. For example, a plate heat exchanger economizer was installed with a 500-ton cooling system at a data center in 1983, with an expected energy savings of 35%. However, due to concerns that valves controlling the water flow through the economizer might get stuck, or that the control sequence might fail, the economizer was never used. [73]
- Variable speed drives are a well-known technology to save energy when moving fluids such as air and water. Because the energy required to move air increases as the cube of fan speed, digital control technologies that adjust HVAC ventilation fan speeds to meet load requirements can improve energy efficiency greatly. Similar control technologies can be employed with chilled water system circulation pumps to reduce energy use.
- Fresh air cooling uses filtered ambient air directly for cooling when outside temperatures are low enough. This eliminates compressor and condenser energy use, though energy may still be required for reheating and humidity control. In appropriate climates this technique has the potential to significantly reduce HVAC energy requirements, though the HVAC system will still have to be sized to meet peak daytime loads.
- Cooling systems can be designed in a modular fashion, such that additional modules become activated as cooling requirements increase, and deactivate as cooling requirements decrease.
- It is estimated that absorption chilling could reduce cooling loads by 30%. [74] However these technologies may currently require more maintenance and be less reliable than other solutions. Some absorption chillers require temperatures of 250 °F or greater, though one water-fired single-effect chiller can utilize hot water at 167 °F - 212 °F. [75] This heat could be generated by distributed generation (DG) technologies such as fuel cells or micro-turbines in combined heat and power (CHP) applications. For this application, distributed generation should be sized to provide enough heat for absorption chilling, but not more, because production of unused waste heat reduces system efficiency.
- According to EPRI, desiccant cooling is applicable in any situation where humidity control is an essential component of indoor temperature control. This is the case in data centers, In standard air conditioning equipment, 33 percent of total energy use is required for humidity control. By incorporating a desiccant system, energy required for humidity control by the air conditioning equipment is reduced to 11 percent. Because latent heat load has been met by the desiccant system, conventional air conditioning equipment size is greatly reduced. All that is left for the air conditioning equipment is to lower the air temperature. Depending on the application, total air conditioning electricity consumption can be reduced up to 40 percent. [76]
- Spray cooling offers a possible future cooling technology for data centers. [77] Spray cooling systems continuously chill water at night in a spray process over large, low-slope roof surfaces. The chilled water is captured at roof

drains and stored for the next day's cooling needs. Spray cooling is applicable in dry climates with clear summer night skies, such as the western U.S. While spray cooling is not a well-known technology, it has potential to increase energy efficiency by taking advantage of free cooling from the night sky, as well as allowing for downsized chiller units. [78]

Air Handling System Improvements

Maintenance of internal server temperatures and relative humidity within operating tolerances is a critical issue for data centers. Current practice is to mount server racks on a raised floor plenum area. Computer room air-conditioning (CRAC) units blow cold air into the plenum, pressurizing the space under the racks and forcing cool air up through perforated floor tiles. Server cooling fans then circulate this air through the rack cabinets, and exhaust it back to the room where air handling units draw off the waste heat.

However, current data centers often have problems with temperature hot spots, and industry experts believe the heat load from future ultradense servers may soon overwhelm conventional cooling technologies. Proper rack layout, energy audits and comprehensive maintenance procedures, and advanced cooling technologies can save energy while improving cooling and reliability.

- Many server farms are laid out quickly and haphazard placement of server racks and poor routing of cables under the plenum floor can impede efficient cooling. The Uptime Institute states that a best practice is to use rows of server racks in alternating "hot" and "cold" aisles. Servers should face cold aisles with perforated floor tiles to access cold air, and should blow waste heat out into hot aisles which have no perforated tiles. [79]
- Stable relative humidity (RH) is important for servers. Decentralized humidifiers in CRAC units can loose calibration, one raising RH while the other lowers it, thereby wasting energy. Efficiency gains can be realized by using central RH control instead. [80]
- Preventative maintenance procedures are important for maintaining cooling system efficiency. The Uptime Institute found that 10 - 30 % of existing cooling capacity is wasted or cannot be realized due to design deficiencies or operational problems. Problems they found at a number of data centers include; dirty or blocked cooling coils, control points located such that they are unable to control HVAC systems accurately, uncalibrated or damaged sensors, reversed supply and return piping, pumping system and valve problems, lack of quality control, repairs partly or improperly done, or not done at all, to name a few. Thus, a simple but rigorous regime of preventative maintenance can save energy at essentially no cost, and may avoid future repairs or downtime. [81]
- The Uptime Institute states that current cooling methods are sufficient to total computer room power densities of 40 W/ft2, but as densities increase conventional cooling technologies will be insufficient. They expect server racks to trend from current rack power densities of 600 W/ft2 to 1600 - 1800 W/ft2 by 2010. Telecommunication frames, which are typically more energy intensive than server racks, may reach product densities of 7000-9000 W/ft2 by 2010. [82] In response to expected increased cooling loads, two different cooling technologies have recently been released:
- On September 10, 2001, RTKL Associates announced a new air-based cooling technology for data centers, termed "high delta T cooling" (HDTC). This system directs cool air from the floor plenum directly through the rack cabinets and into a return plenum, so cooling air is always focused right on the servers. Capturing the waste heat directly doubles the temperature difference across the cooling coil compared to standard systems, doubling cooling efficiency. Because HDTC uses half the airflow of conventional systems, the number of CRACs is cut by half, lowering cooling capital and installation costs. RTKL claims that HDTC reduces overall load by 6 - 10% while reducing data center cooling costs by up to 14%. Trial

installations have been successful at two different Fortune 500-owned data centers. [83]

- Also on September 10, 2001, Sanmina Enclosure Systems announced a new water-based cooling technology for data centers, the "Ecobay 442". The system is an is an assembly of four enclosed cabinets with an integral cooling module and power and control electronics capable of housing 160 servers. It utilizes a closed loop chilled water system running throughout the cabinet to draw off heat. As with the RTKL technology, cooling the servers directly allows ambient room temperatures to be set higher, further reducing conventional HVAC loads in the data center. Sanmina claims cooling energy savings can reach 50 percent in some cases, and optimizing energy efficiency in ultra high-density data processing applications. [84]

Balance-of-System Component Improvements

Aside from HVAC and server power demands, balance of system components such as uninterruptible power supplies (UPSs) and power distribution units (PDUs) use energy. While these systems are 95-98% efficient under full design load, as we have seen data center loads are much less than design loads. It is estimated that 5-7% of the incoming power is lost as it passes through UPSs and an additional 2-5% of the remaining power is lost to PDUs. Electrical line losses within the data center also account for approximately 1% of incoming power. [85] Sizing UPSs and PDUs to meet actual loads can improve energy efficiency, and additional R&D may help reduce these losses further.

Future Prospects for Improvements

A number of efforts are underway to further improve the energy efficiency of the digital information sector. For example, the U.S. Defense Advanced Research Projects Agency (DARPA) recently awarded a $2 million grant to Northwestern University to develop software and power-aware processor architecture that will work with one-another to improve energy efficiency. The goal of Northwestern's Power Aware Architectural and Compilation Techniques (PACT) project is to reduce power demand by factors of 10 to 100. [86] Also, the NASA Ames Research Center is conducting research on use of optics for data transfer and optical switching. These new technologies that may significantly increase computing speed and reliability while reducing energy requirements. [87]

EFFICIENCY TARGETS

Based on the range of potential energy savings identified above, we estimate impacts on energy demand by end-use at different levels of efficiency implementation to arrive at plausible efficiency targets. We build on measured end-use data to estimate the effect of these efficiency targets on data center components by end use, and aggregate these effects to estimate impacts on a generic "model" data center. We then use this result in combination with industry data center growth projections to estimate potential national energy savings attributable to data center efficiency target implementation.

Efficiency Targets for a Model Data Center

Measured energy data by end-use for one data center is given in Table 6. [88] While this does not necessarily represent the end-use profile for all data centers in the U.S., it provides a basis for exploring the impact of energy efficiency and smart design measures described in the previous section. Based on historical trends, recent technology developments, and interviews with industry experts, we estimate reasonable targets for energy efficiency improvements by end use. In the first estimate, we assume only efficiency increases that are realizable through minimal effort. In the second estimate, we assume an aggressive efficiency regime, taking advantage of the best available technology on the market, as well as technological advances expected in the near term.

Based on available data, we assume generic data center owners will continue to request enough power to provide a total computer room power density of at least 100 W/ft2. Because reports indicate data centers consume approximately a third of their original requests, we assume actual power densities to be 35 W/ft2. With minimal energy efficiency improvements in all end-uses, we assume

Table 6. Total Computer Room Peak Demand by End-Use: Nameplate, Measured, and Efficiency Targets [*]

End-Use	Measured contribution to load (%) [a]	Nameplate-based demand estimate (W/ft2) [b]	Measured demand estimate (W/ft2) [c]	Minimal efficiency target (% reduction) [d]	Minimal efficiency target (W/ft2)	More aggressive efficiency target (% reduction) [e]	More aggressive efficiency target (W/ft2)	Potential Efficiency Improvements [n]
computers	49%	48.6	17.0	25%	12.8	65%	6.0	low-wattage processors [h]
air handling (fans, CRAC units, AHUs)	23%	23.2	8.1	20%	6.5	50%	4.1	direct high delta-T air-cooled or water cooled racks, variable speed fans, improved preventative maintenance [i]
central chiller plant [f]	14%	13.6	4.8	10%	4.3	30%	3.3	state-of-the-art chillers, economizer cycles, outside-air cooling, absorption chilling, desiccant cooling, spray cooling [j]
auxiliary equipment (UPS, PDU), other	11%	11.1	3.9	5%	3.7	20%	3.1	advanced circuit design [k]
lights	3%	3.4	1.2	50%	0.6	75%	0.3	high-efficiency T-5 and T-8 fluorescent with magnetic ballasts and automatic dimmers [m]
Total peak computer room power density (W/ft2) [g]		100	35	20%	28	52%	17	

[*] Please see Appendix A for notes, references, and explanation of assumptions for Table 6.

this power density could be reduced by 20% overall to 28 W/ft2. Going further, with aggressive energy efficiency measures in all end-uses, we assume data center power requirements could be cut by more than half (52%) overall, reducing power density requirements to 17 W/ft2. The result is that power density requests based on actual energy use and minimal efficiency improvements could provide a four-fold decrease in power requirements from an original nameplate based request (from 100 W/ft2 to 28 W/ft2), and aggressive efficiency improvements lead to a six-fold drop from that original request (from 100 W/ft2 to 17 W/ft2). While power densities will vary from data center to data center, this example provides an order-of-magnitude estimate of what could be achieved on the level of individual data centers through energy efficiency improvements.

National Efficiency Targets for Model Data Centers

The impact of these energy efficiency targets can be generalized to estimate energy savings impacts on data centers nationally. These estimates and impacts are presented in Table 7. As a baseline for gauging national impacts, a national laboratory study determined that data centers by themselves consumed at most 0.15 to 0.2 percent of U.S. electric energy as of the end of 2000. [89] This translates to 5,100 to 6,800 gigawatt-hours of consumption in 2000, requiring between 580 and 780 MW of continuous supply on average. [90] Based on published data center square footage surveys and energy requirements of 35 W/ft2, we estimate actual U.S. data center peak demand in 2000 was 333 MW. This is approximately half of the national laboratory estimate for that year. However, the data center survey used for our estimate did not account for all data centers in the U.S., which may account for some of this discrepancy.

In August, 2000, Salomon Smith Barney estimated that total U.S. data center computer room floor space would increase from approximately 9.5 million square feet in 2000 to 25 million square feet in 2003. Due to the recent economic downturn this growth in data center floor space will likely be delayed, but the following serves to illustrate the energy implications of data center growth as the economy picks back up. Applying the power density assumptions from Table 6, we estimate total U.S. data center power demand in 2000, 2001, and 2003 at both nameplate-based and measured use levels. These estimates are based on projections of raised floor area. To arrive at industry-wide impacts of efficiency improvements, we assume the end-use efficiency improvements derived in Table 6 are implemented in all data centers newly built or expanded after 2000. We further assume that, on average, the energy savings shown in Table 6 are realized on a national scale.

Based on past trends and available data, we assume data center owners can reasonably be expected to demand at least an additional 1,550 MW as data center construction or expansion increases raised floor area nationwide from 9.5 million to 25 million ft2. Because reports indicate data centers draw only a third of their original request, we expect only about 540 MW of new peak demand to be realized. With minimal energy efficiency improvements this demand could be reduced by 20% (110 MW) to 430 MW. Going further, aggressive energy efficiency measures could cut data center peak energy demands by more than half, reducing demand by 280 MW for a final new demand of 260 MW. Thus, if data center owners aggressively employ energy efficient technology and smart design, and base power requests on actual and not over-estimated loads, peak power demand requests of utilities from data centers could be dropped six-fold, from 1550 MW to 260 MW. A milder efficiency effort combined with measured-use demand estimates would still drop new peak demand requests fourfold, from 1550 MW to 430 MW.

A May 2001 article estimated that between 5,000 MW and 10,000 MW of new electricity demand from data centers will be realized in the U.S. by 2005 from data centers requests currently in the pipeline. [91] This is the equivalent of 17 to 33 three-hundred megawatt power plants. Should this demand be realized, data center energy savings at either of the efficiency target levels defined above become significant on a national level, avoiding construction of three to seventeen new power plants over the

Table 7. Estimated U.S. Data Center Peak Demand Reductions From Efficiency Improvements, Based on 2000 Growth Projections, 2000-2003 **

Year	Estimated computer room space (ft2) [o]	Energy demand at 100 W/ft2 nameplate (MW)	Energy demand at 35 W/ft2 measured (MW)	Energy demand with minimal efficiency improvements (MW)	Energy demand with aggressive efficiency improvements (MW)
2000	9,500,000	950	333	265	160
2001	17,900,000	1790	627	499	301
2003	25,000,000	2500	875	698	420
growth 2000-2001	8,400,000	840	294	234	141
growth 2000-2003	15,500,000	1550	543	432	260
efficiency savings potential from new construction, 2000-2001 (MW) [p]				60	153
efficiency savings potential from new construction, 2000-2003 (MW) [q]				110	282

next four years. Because data center development is currently stalled due to economic conditions, this level of demand will likely not be realized in 2005, but sometime after that.

In addition to efficiency improvements in specific data center end-uses, further efficiency and/or reliability gains may be possible through onsite (or "distributed") power generation, as well as through co-locating multiple data centers in a "power park" setting, allowing them to take advantage of synergistic efficiency measures such as district cooling that an isolated data center could not take advantage of. While a full treatment of these approaches to data center efficiency is beyond the scope of this paper, Box 4 provides a brief overview of these concepts.

NEXT STEPS

With available energy efficient technologies and identified energy-saving design and maintenance procedures, what will it take to reduce data center energy demand toward our proposed efficiency targets while maintaining or improving reliability? This section explores a few of the possible next steps to move utility and data center industries toward this goal.

** Please see Appendix A for notes, references, and explanation of assumptions for Table 7. The years 2001 and 2003 are given for reference only. Energy demand will be driven by actual raised floor area expansion, rather than the year it occurs in. Given the dot.com crash, 2001 economic slowdown, and economic impact of September 11, 2001 events, it is almost certain that these raised floor space estimates will not be realized in the year specified in the pre-downturn data center projections. However, as the economy recovers data center floor space will once again grow. As it does, these estimates of efficiency savings can be applied accordingly.

Box 4: Onsite Generation and Power Park Concepts

Distributed Generation and Combined Heat, Power, and Cooling
As more energy-intensive data centers come on line at a time when electricity supply is tight in many areas of the U.S., customer-sited power, or distributed generation (DG) is increasingly being considered as an option to provide some or all of the data centers power and cooling needs. When U.S. DataPort announced plans to build a 2.2 million square foot facility as part of a 174-acre project in San Jose, PG&E warned it might not be able to supply the projected 180 MW load. The project was approved in April 2000, but U.S. DataPort had to agree to build a 250 MW generating plant to power the project. [92]

Debate over the pros and cons of distributed generation is of ongoing interest as utilities continue to deregulate. DG is most efficient in combined heat, power, and cooling (CHPC) applications, where waste heat from power generation can be used to power absorption chillers to cool the data center. It has been estimated that CHPC could save 30% of overall load that chillers would have used. [93] However, siting of large (50 MW and up) DG power plants in urban areas have proved to be controversial from an environmental and public standpoint. [94]

Renewable Distributed Generation with Photovoltaics
In one success story for clean distributed generation, Internet hosting company SolarHost has been running a small web hosting facility partially powered by solar photovoltaics in northern Virginia since March 2000 with only 11 minutes of downtime since inception, mostly for service upgrades.[95] SolarHost is considering installing a solar-powered data center in Tampa with 2000 ft2 of raised floor space that will eventually house 250 servers. The data center and commercial space, 6000 ft2 overall, would be entirely powered by 1.2 MW of solar photovoltaics on the buildings roof. The solar panels would be used to charge a bank of batteries that provide a full five days of backup power. Future expansion is possible, as the 40,000 ft2 of available roof space can house up to 4.2 MW of photovoltaics. While SolarHost services cost 20 – 25% more than industry average, there is enough interest in clean, renewably-powered Internet hosting that they have been financially successful. [96] Companies that want to meet corporate environmental commitments can look to SolarHost's example to see that web hosting powered by renewable distributed generation is marketable in the U.S. and may help them meet their social goals.

Smart Power Parks
Because DG has significant upfront capital costs, may incur opposition from environmentalists, and isolated high loads create risk exposure when faced with volatile customer bases, data center owners may be hesitant to pursue DG options. The "smart power park" concept--building a number of data center facilities in one location, powered by a dedicated power plant--may address some of these concerns. By having multiple customers, risk of stranded generating investments is minimized. Creating a mixed use environment, with other industrial or commercial customer relying on the dedicated power plant, would further serve to reduce risk. Use of DC-power microgrids, rather than AC power, may increase overall energy efficiency. [97] With CHP applications, waste heat could power absorption chilling which could then be distributed throughout the power park using a district cooling loop. In this way all waste heat could be put to use, whereas powering an individual data center solely with DG would likely result in excess waste heat and reduced efficiencies. District cooling is an established technology and is being used in cities across the country in large buildings and developments. Austin Energy has a number of district cooling projects in place and underway, and has considered using it in the power park scenario. [98]

The U.S. Department of Energy Office of Power Technologies has programs related to both distributed energy resources and power park applications. More information on these programs can be found at the Northeast-Midwest Institute Distributed Energy Resources (DER) website, <http://www.nemw.org/energy_der.htm >.

Rewarding Efficiency

Data center operators have significant disincentives to making their data centers energy smart. Facility managers responsible for running data centers receive no mandate from higher management to save energy in data center operation. They are heavily penalized for any service interruption or perceived decrease in reliability, and not rewarded for saving energy through efficient technology or better maintenance practices. This could be changed if higher management and executive levels of the companies were to realize the value of energy efficiency in improving the bottom line of their company.

One example of involving heads of major corporations in energy-saving programs is the Climate Savers Program, a joint project by the World Wildlife Fund and the Center for Energy and Climate Solutions. This program helps businesses reduce climate change impacts from energy consumption by adopting energy efficient and clean energy technologies in both commercial and industrial settings. [99] Program members currently include IBM, Johnson & Johnson, Polaroid, and Nike.

Another example is "Environmental Charter for the North American Telecommunications Industry." This voluntary environmental charter identifies telecommunications industry areas of environmental responsibility and seeks to minimize its environmental impacts (among other goals). A number of major telecommunications companies have signed onto the charter, including Ameritech, AT&T, BCT.Telus, Bell Atlantic, Bell Canada, Bell South, and US West. [100] Programs like Climate Savers and the telecommunications environmental charter targeted to the ICT community could help parent companies of data centers meet social and environmental goals through energy efficiency while achieving positive public and media attention for their commitments.

Finally, Savings by Design, a unique program funded by California utility customers and run by California utilities, helps architects and building owners increase the energy efficiency and productivity of commercial buildings. They offer financial incentives to building owners of up to $150,000 per building, and to design teams of up to $50,000 per project for improvements in building energy efficiency. These incentives are in addition to the value of avoided energy and capital costs from efficient building design. This model could serve as a template for a similar utility program to incentivize data center energy efficiency. [101]

Energy Service Companies

Another way to incentivize energy savings is to turn the management of energy issues over to independent contractors known as Energy Service Companies, or ESCOs. These companies are firmly established in providing energy management services and energy efficient improvements in the commercial buildings sector. Because ESCOs derive their revenue from avoided energy charges, they have the incentive both to save energy through increased energy efficiency and to retain their customers in the process. However, security and reliability concerns make data center owners and managers extremely skeptical about the prospect of turning over control of their power supply and cooling management to a third party such as an ESCO. These concerns need to be explored, defined, and addressed before ESCOs can be expected to make a difference in the data center sector. [102]

Industry Cooperation

As with other economic sectors, cooperation between data center industry actors can go a long way toward eliminating the inefficiencies of piecemeal approaches to design solutions. The U.S. Department of Energy Buildings Program has called for a reduction in the fragmentation of the buildings industry and a sharing of knowledge between actors in designing efficient buildings. [103] This type of effort is applicable to the data center industry, and is already underway. For example, in February, 2001, Edison Electric Institute (EEI) held a seminar to bring together utilities and data center industry actors to address the need to provide reliable power to data centers without putting utility infrastructure investments at risk. The need for developers, design engineers, utilities, telecom providers, and others to work together was identified, and an Internet

Infrastructure Consortium was created to provide a forum for making progress on this issue. [104] The EEI consortium held its second data center conference in September 2001 and is planning a third conference in Spring, 2002. The consortium is currently working to develop a guidebook for data center developers to broaden their understanding of power supply and demand. [105]

Industry Design Standards

An example of industry cooperation would be the development of voluntary industry design standards for best efficiency practices in data center design and technology. A number of industry organizations are already working toward this goal. EPRI's Consortium for Electric Infrastructure to Support a Digital Society (CEIDS) program has a multi-million dollar R&D budget earmarked to solve critical infrastructure challenges in the digital society. [106] The Uptime Institute, a group of Fortune 500 companies working to improve reliability management in data center facilities and information technology organizations, has published a number of white papers on site reliability procedures and guidelines, and has recently released a draft "Cooling Compatibility Specification" for industry comment. [107] The PCI Industrial Computer Manufacturers' Group (PICMG) is a consortium of over 700 companies that collaboratively develops open specifications for high performance telecommunications and industrial computing applications. [108] The 7x24 Exchange is an industry forum for improving "End-to-End" reliability of mission-critical enterprises such as data centers. [109] These organizations and others like them can play an important role in developing voluntary standards for energy saving architecture and design.

Rating, Certification, and Labeling Programs

Rating, certification, and labeling programs have helped energy efficient technologies gain a toehold in the market. For example, the U.S. Environmental Protection Agency 's (EPA) ENERGY STAR [110] computer program helped promote the transfer of energy-saving features of laptop computers, such as automatic screen savers and standby and sleep modes, to desktop computers. A 1993 Executive Order required that microcomputers purchased by federal agencies be ENERGY STAR compliant. This prompted computer manufacturers to participate in the program. As of 2000, EPA estimates that 80% of computers sold in the U.S. are ENERGY STAR compliant. [111]

How can this benefit the data center industry? Building on past work, the ENERGY STAR program could potentially be used to accredit "Energy Smart" servers and/or "Energy Smart" data centers. EPA's recent ENERGY STAR "certificate of recognition" awards for IBM's eServer z900 and Transmeta's Crusoe microprocessor indicate that EPA may be ready to consider rating servers. Similarly, EPA's ENERGY STAR buildings program could be extended to cover data centers. This program has established benchmarking tools for rating the energy efficiency of different building types, including offices, schools, and hospitals, and comparing them to their peer group. Buildings in the top 25% tier of most efficient buildings may be qualified to receive an ENERGY STAR building label.

As it turns out, this effort is already underway. EPA's ENERGY STAR building program is developing a benchmarking tool for data centers and telecommunications equipment buildings. (Telecommunications centers are similar to data centers but are generally more energy intensive.) Because Verizon is the only major telecommunications company with an accurate energy data management system, they have partnered with EPA and have agreed to provide their telecommunications center energy use data to EPA. (Verizon consumes over 5 billion kWh of electricity annually.)This data will be used to calibrate EPA's new data center benchmarking system. [112]

Additional roles for ENERGY STAR rating programs may apply to markets for more efficient power supplies and more efficient computer memory. [113] As discussed elsewhere in this paper, power supplies may be a good candidate for additional energy efficiency improvements. And, as the need for data storage grows, reducing energy use by computer memory may become more important. Entities other than EPA could be formed to rate,

accredit, and label energy efficient data centers or their components. For example, in the windows industry the National Fenestration Rating Council (NFRC) was formed to voluntarily rate and label the energy performance of windows. This was done to avoid having the Federal government imposing a rating and labeling system itself.

Load Characterization

As illustrated by EPA's need to acquire power consumption data from Verizon, measuring, tracking, reporting and characterizing data center loads accurately is important for the data center industry, as well as for utilities. Good data allows accurate utility load planning, identification of areas for data center efficiency improvements, as well as verification of the impact of efficiency improvements. At the February, 2001, Edison Electric Institute (EEI) data center industry meeting the need for accurate data for ultra-high density loads such as data centers was clearly identified as a priority, and in June, 2001, EEI launched a "Internet Hotel" data survey to begin to quantify data center loads. [114] Lawrence Berkeley National Laboratory (LBNL) will also be collecting energy use data for several data centers under funding from the New York State Energy Research and Development Authority (NYSERDA) and the California Energy Commission. [115] However, data collection may be both expensive and inexact because many data centers have multiple tenants per building (50-100 for some co-located facilities) and short (1-3 year) leases. It could be very difficult to document and track loads in a useful manner because IT technology changes rapidly and customers move in and out often. In addition, any information regarding the operation of data centers is usually very tightly guarded for competitive and security reasons. Establishment and publication of data center energy demand surveys such as EEI's and LBNL's would be a great help in load characterization.

Such an effort could be undertaken by the U.S. DOE's Energy Information Agency (EIA). The EIA already has in place a Commercial Buildings Energy Consumption Survey (CBECS). The CBECS is a national sample survey began in 1979 that collects energy-related building characteristics data and energy consumption and expenditures data for commercial buildings in the United States. While EIA is not planning a specific study on data centers, they will be adding additional questions on computer and internet usage to their upcoming 2003 survey. [116]

Balance-of-System Component Improvements

In addition to improving processor and HVAC efficiencies, much more can be done to reduce the power use of balance-of-system components such as UPSs, PDUs, and other digital electronics. Out of a $231 million proposed budget for CEIDS, EPRI has earmarked $10 million to develop advanced circuit topologies for improving switch-mode power supply performance, thereby reducing total harmonic distortion (THD), increasing energy efficiency, providing 1-2 seconds of ride through capability (60 to 120 times longer than standard power supplies), and providing improved surge/transient immunity to reduce the impact of power outage and power quality issues. [117]

Zoning Codes

Zoning classifications affect HVAC requirements and thus energy use. Data centers zoned as office buildings require a much higher number of air changes per hour (ACH) than data center zoned as warehouses because zoning codes assume office buildings require more fresh air for their workers than warehouses. Thus, even though two data centers might be identical, the one zoned as an office will use more energy due to increased HVAC activity. Because data centers are a relatively new phenomenon, and are rapidly changing in configuration and location, many zoning laws do not treat data centers as a specific sector. Amending existing commercial zoning laws to take into account unique characteristics of data centers may pave the way for HVAC-related reductions in data center energy demand.

CONCLUSION

Reasonable assumptions, based on historical and current trends, lead us to set efficiency targets of a 20% reduction in data center total computer room measured power density with a minimal efficiency effort, and a 52% reduction in measured power density with an aggressive efficiency effort. Were these improvements implemented industry-wide, analogous reductions in peak demand could be realized on a national scale.

We find the combination of suggestions detailed in this paper—basing data center energy requests on measured loads rather than nameplate-based loads, and employing energy-efficient technology, energy smart design guidelines, and improved maintenance routines—could reduce data center peak demand requests fourfold from current levels in the minimum efficiency case, and six fold in the aggressive efficiency case. These reduced demands should be much easier for utilities to meet, as compared to the high demand projections touted by the media prior to the recent economic downturn.

While technological and design improvements to reduce power demands of specific data center end-uses are important, the greatest efficiency gains will come when data centers are considered as a system. As shown by tried and true efficient residential, commercial, and industrial facility design, exploiting synergies between different end use technologies and design improvements will provide savings greater, often much greater, than will a piecemeal approach. Similarly, cooperation between all actors involved in the data center field, from technology suppliers, designers, contractors, and facility operators to utility suppliers, can maximize data center energy efficiency while maintaining or improving reliability.

Better understanding of the issues discussed in this paper can help further improve digital technology energy efficiency. These improvements may benefit not only data centers, but a number of commercial and industrial processes and applications that rely on digital technologies.

APPENDIX A: FOOTNOTES AND ASSUMPTIONS FOR TABLES 6 AND 7

[a] These values are based on January 2001 measured energy use at one SF Bay Area data center as reported in "Table 8. Breakdown of Power Density by End Use" by Jennifer D. Mitchell-Jackson. "Energy Needs in an Internet Economy: A Closer Look at Data Centers" May 16, 2001. Master's Thesis, Energy and Resources Group, University of California. Berkeley, California. Available at <http://enduse.lbl.gov/Projects/InfoTech.html>

PG&E has found slightly larger server and air handling loads, at 55% and 32% of load, respectively. Personal communication with Keith Reid, PG&E Senior Corporate Account Manager, November 8, 2001.

[b] A recent survey by a major IT website indicated that 85% of data center operators polled maintain power capacity to feed a total computer room power density of at least 100 W/ft2, and over 20% of all operators polled maintain capacity to feed power densities of 200 W/ft2 or more.

We make the conservative assumption that the average data center designs to 100 W/ft2. Survey data published in Rich Miller, CarrierHotels.com. "How Much Power? Survey Shows Split Among Facility Operators." (March 21, 2001). Accessed August 15, 2001 at <http://www.carrierhotels.com /wiredspace/wired0321.shtml>

[c] Measurements of data center power use indicate that actual loads are one-fifth to one-third of the requested design load. For example, PG&E has reported that a fully occupied data center in their territory drew one-third of the requested amount. To be on the conservative side, we assume actual loads are 35% of design loads. This is supported by personal conversations with Keith Reid, PG&E Senior Corporate Account Manager, July 11, 2001 and Chris Wilkins, P.E., Hallam Associates, August 1, 2001; as well as Stuart M. Lewis. Transmission & Distribution World. "Utilities Cannot Afford to Become 'Sometimes Power & Light'" (April 2001); and Jennifer D. Mitchell-Jackson. "Energy Needs in an Internet Economy: A Closer Look at Data Centers" May 16, 2001.

[d] These are estimates on the lower end of what power use reductions may be possible using current technologies listed for each category. Where possible these efficiency targets are based on available literature, but in some cases are speculative and have not been verified by our research.

[e] These are estimates on the higher end of what power use reductions may be possible using best available technologies on the market today, as well as improvements that may be made in the near term. Where possible these efficiency targets are based on available literature, but in some cases are speculative and have not been verified by our research.

[f] As measured by Jennifer Mitchell-Jackson, combined use of chiller and air conditioning and air handling equipment accounts for 37% of the data center power density. This is in good agreement with other industry sources, which estimate total HVAC energy use at 30-40% of total data center power density. Russ Zabel. South County Journal. "`Server Farm' Boom Poses Threat to NW Power Supply" (Seattle: December 26, 2000); Thomas Callsen. "Internet Hotels--A Big Draw" Transmission & Distribution World. February 2001; and personal conversation with Keith Reid, PG&E Senior Corporate Account Manager, July 11, 2001.

[g] Figures quoted in percent (%) denote the percent decrease in power use from measured loads values from efficiency improvements.

[h] From historical data, it is known that development of CMOS-based mainframes in the mid-1990s replaced previous liquid-cooled mainframes, cutting energy use by 95%. Recent processor developments may reduce energy use of existing Pentium chips from 33-50 watts per chip to half a watt per chip, a drop of two orders of magnitude. Further, IBM expects that new processors may use one-tenth the energy of current processors. Finally, NetWinder states their Crusoe-based server will use one-quarter the energy of comparable Intel-based servers. Thomas Callsen. "Internet Hotels--A Big Draw" Transmission & Distribution World. (February 2001); Jonathan Angel. Network Magazine. "Emerging Technology: Energy Consumption and the New Economy" (January 5, 2001); IBM Press release. "IBM Launches 'High-Powered' Initiative for 'Low-Power' Products and Services" (Armonk, NY; October 1, 2001); IBM Press release. "IBM Unveils Revolutionary Low-Power Chip Technologies" (East Fishkill, NY: October 12, 2001); and George A. Chidi Jr., IDG News Service. "New servers use Transmeta chips to cut power costs" (Boston: January 25, 2001) <www.idg.net>.

[i] The Uptime Institute found that 10-30% of cooling capacity is wasted simply due to poor design and maintenance. RTKL's new forced-air cooling technology cuts the number of CRACs by 50%, reducing overall cooling energy by 14%. Sanmina claims their new chilled-water server racks can produce overall energy savings of up to 50%. Higher estimates of HVAC energy savings assume more efficient servers are used, reducing overall heat loads. W. Pitt Turner IV and Edward C. Koplin. Uptime Institute white paper. "Changing Cooling Requirements Leave Many Data Centers at Risk" (Uptime: 2000); R. Stephen Spinazzola, P.E., and Zubin Menachery, P.E.. RTKL Associates white paper. "High Delta T Cooling (HDTC) provides increased energy efficiency and reliability for data centers." (Baltimore: 2001); and Sanmina Enclosure Systems. Press release. "New Sanmina Product Dramatically Reduces Energy Costs in Data Centers" (Toronto: September 10, 2001).

[j] RTKL claims their new forced-air cooling technology reduces overall cooling plant energy by 14%. It has been estimated that absorption chilling in combined heat and power applications could save 30% of chiller load. According to EPRI, desiccant cooling can reduce total air conditioning energy consumption up to 40%. R. Stephen Spinazzola, P.E., and Zubin Menachery, P.E.. RTKL Associates white paper. "High Delta T Cooling (HDTC) provides increased energy efficiency and reliability for data centers" (Baltimore: 2001); Personal conversation with , P.E., Hallam Associates, August 1, 2001; and EPRI. Consortium for Electric Infrastructure to Support a Digital Society (CEIDS) "Draft R&D Plan" (March 7, 2001).

[k] Efficiency gains of 5% are possible simply by running UPSs and PDUs at full rather than part load. Additional gains can be made from improved circuitry that improves power quality. Thomas Callsen. "Internet Hotels--A Big Draw" Transmission & Distribution World. (February 2001); Jennifer D. Mitchell-Jackson. "Energy Needs in an Internet Economy: A Closer Look at Data Centers" May 16, 2001; and EPRI. Consortium for Electric Infrastructure to Support a Digital Society (CEIDS) "Draft R&D Plan" (March 7, 2001)

[m] Commercial buildings have saved 50% of lighting energy through basic upgrades. An engineering study performed for SolarHost found energy efficient lighting reduced data center lighting energy density by 75%. Boeing reduced lighting energy use by 90% in a several million square foot facility with a 50% return on investment. Joe Romm. Center for Energy and Climate Solutions. "Cool Companies" (Washington DC: 1999); and personal communication with Steven May, President, SolarHost. (October 5, 2001).

[n] While suggested efficiency targets are proposed for individual data center subsystems, it is critical to realize that all subsystems interact with one-another. Maximum efficiency gains can only be realized when synergies between system technology, design, and maintenance improvements are explicitly taken into account.

[o] This estimate is based on surveys of data center computer rooms built or under construction conducted by Salomon Smith Barney (Mahedy, Stephen and Dan Cummins and Danny Joe. "Internet Data Centers: If Built...Will They Come," Salomon Smith Barney Report, 3 August 2000), the Yankee Group ("Executive Summary of The U.S. Collocation Market: High-Tech Real Estate Heats Up," 2000, <www.yankeegroup.com>), and Robertson Stephens (Juarez, Richard A and Michael T. Alic, Chetan S. Karkhaniz and Brett D. Johnson. SpaceDexIII. Hosting Space: Not All Space Is Created Equal—Smart, Complex Space Takes Center Stage," Robertson Stephens, 29 January 2001) as adapted by Jennifer D. Mitchell-Jackson in "Energy Needs in an Internet Economy: A Closer Look at Data Centers" (May 16, 2001). While economic conditions may result in less data center computer room space being built by the end of 2001 and 2003, this table provides an estimate of the magnitude of energy savings possible through energy efficiency measures.

[p] Based on efficiency assumptions as noted in Table 6.

[q] Based on efficiency assumptions as noted in Table 6.

ENDNOTES

[1] "Data center" will be used throughout as a generic term to describe a range of facilities that use high concentrations of computers, servers, and digital electronic equipment to host websites, store, process and transfer electronic data, and otherwise support a digital electronic economy.

[2] American Gas Cooling Center. *Cool Times Energy Solutions.* " Natural Gas Technologies Promise to Feed Power-Hungry Data Centers" (Spring 2001).

[3] David Kathan and Thomas J. Grahame. *Broadband Wireless Online.* "Internet Data Centers: Demands for Electricity Proceed Unabated. Part 2 of 2" (June/July 2001). Vol. 2, No. 6. Accessed October 18, 2001 at <http://www.shorecliffcommunications.com/magazine /volume.asp?Vol=16&story=164>

[4] According to a New York Times article, today there are six million servers in the world, while in 1995 there were just 20,000. Roger Anderson. "Wattage Where It's Needed." New York Times Op-Ed. (New York: June 6, 2001).

[5] The Emerging Digital Economy II. U.S. Department of Commerce, Economics And Statistics Administration, Office of Policy Development, Washington, DC. ESA/OPD 99-6. June 1999. Accessed February 12, 2001 at <http://www.esa.doc.gov/508/esa/TheEmergingDigitalEconomyII.htm>

[6] Digital Economy 2000. U.S. Department of Commerce, Economics And Statistics Administration, Office of Policy Development, Washington, DC. June, 2000. Accessed February 12, 2001 at <http://www.esa.doc.gov/508/esa/DigitalEconomy.htm>

[7] Iwata, Edward, "Online Revolution Here to Stay: Venture Capitalists Say Internet is Under-Hyped," San Francisco Examiner, June 10, 1999 and Forrester Research, Inc., 1999 as cited in John A. "Skip" Laitner. The Information and Communication Technology Revolution: Can It Be Good for Both the Economy and the Climate? [Discussion Draft]. EPA Office of Atmospheric Programs. Washington, DC. December 15, 1999. (Revised January 31, 2000) Accessed February 12, 2001 at <http://enduse.lbl.gov/projects/ITrevolution.pdf>

[8] SPS/Spectrum Economics. "Global Economic and Information Technology Market Forecasts, 2000-2005." September 2000. As cited in Laitner, John A. "Skip," Koomey, J.G., Worrell, E., and E. Gumerman. U.S. EPA and Lawrence Berkeley National Laboratory. "Re-Estimating the Annual Energy Outlook 2000 Forecast Using Updated Assumptions about the Information Economy." Presented at the American Economics Association. (New Orleans: January 7, 2001). Also available as LBNL-46418.

[9] Power quality issues include voltage spikes, surges, sags, line noise, brownouts, blackouts, frequency variations, switching transients, and harmonic distortion.

[10] Joel N. Gordes. Power to Insure: Reducing Insurance Claims with New Electricity Options. September 2000. Northeast Sustainable Energy Association.

[11] EPRI. "The Cost of Power Disturbances to Industrial & Digital Economy Companies" Executive Summary. Prepared for EPRI by Primen. (EPRI: June 2001). Accessed July 2000 at <http://ceids.epri.com/ceids/Docs/outage_study.pdf>

[12] As defined by EPRI, the Digital Economy firms include data centers, telecommunications, biotechnology, electronics manufacturing, and the financial industry.

[13] Draft Environmental Impact Report for the U.S. DataPort Planned Development Rezoning and Prezoning. File No. PDCSH00-06-048. November 2000. City of San Jose, California. Accessed August 15, 2001 at <http://www.ci.san-jose.ca.us/planning/sjplan/eir/USDataport/US-Dataport-Text.htm >

[14] Uptime Institute white paper. "Heat Density trends in Data Processing, Computer Systems, and Telecommunications Equipment. (Uptime, 2000). Available at <www.uptimeinstitute.org>

[15] Mahedy, S., Cummins, D., and D. Joe. Salomon Smith Barney Report. "Internet Data Centers: If Built…Will They Come." (August 3, 2000). As quoted by Tadpole.Com "Tadpole and Platform Computing Launch 'PowerBack' To Reduce Energy Costs For Power-Hungry Data Centers" Press release. (San Jose, CA : September 25, 2001) Accessed October 3, 2001 at <http://www.tadpole.com/cycle/newsevents/press/2001/092501.htm >

[16] Personal communication with Joe Romm, Center for Energy and Climate Solutions. November 5, 2001.

[17] Four-nines of power availability translates to 53 minutes of downtime per year, five-nines translates to 316 seconds downtime per year, and six-nines translates to 32 seconds of downtime per year.

[18] W. Pitt Turner and Kenneth G. Brill. Uptime Institute. "Industry Standard Tier Classifications Define Site Infrastructure Performance." (Uptime: 2001). Available at <www.uptime.com>

[19] Mark P. Mills and Peter W. Huber. Dig more coal--the PCs are coming. Forbes Magazine. May 31, 1999. Accessed at <http://forbes.com/forbes/99/0531/6311070a.htm>. See also: Mills, Mark P. The Internet Begins with Coal: A Preliminary Exploration of the Impact of the Internet on Electricity Consumption. Arlington, VA: The Greening Earth Society. May, 1999. Accessed at <http://www.fossilfuels.org>

[20] Op. Cit. Laitner et. al. LBNL-46418

[21] Center for Climate Solutions. "The Internet Itself Is Not A Big Energy User" Press release. (CECS: December 9, 1999). Accessed October 5, 2001 at <http://www.cool-companies.org/energy/cecs.cfm>

[22] Op. Cit. Koomey et. al, LBNL-44698 and Kaoru Kawamoto, Jonathan G. Koomey, Bruce Nordman, Richard E. Brown, Mary Ann Piette, and Alan Meier. Electricity Used by Office Equipment and Network Equipment in the U.S. Proceedings of the 2000 ACEEE Summer Study on Energy Efficiency in Buildings. Asilomar, CA. August 2000. Also published by Lawrence Berkeley National Laboratory as LBNL-45917. Accessed at <http://enduse.lbl.gov/Info/LBNL-45917.pdf>.

[23] Jonathan Koomey, Kaoru Kawamoto, Bruce Nordman, Mary Ann Piette , and Richard E. Brown. "Initial comments on 'The Internet Begins with Coal'" Memorandum LBNL-44698 (LBNL: December, 1999) . Accessed at <http://enduse.lbl.gov/Projects/InfoTech.html>

[24] Based on work by Jennifer D. Mitchell-Jackson. Energy Needs in an Internet Economy: A Closer Look at Data Centers. May 16, 2001. Master's Thesis, Energy and Resources Group, University of California. Berkeley, California. Available at <http://enduse.lbl.gov/Projects/InfoTech.html>

[25] Watt per square foot (W/ft2) metrics can be misleading when comparing the energy efficiency of data centers because not all such facilities have identical computing power per square foot. Definition of an input/output metric, such as digital information transferred/processed/stored per kWh of power used, could be useful for comparing the energy efficiency of data centers, or at least the servers that populate them. Engineers sometimes use million instructions per second per watt (MIPS/W) as a figure of merit for processors intended for mobile applications. A similar metric is found in the lighting industry, where light bulb efficiency is measured in lumens per watt. Efficient light bulbs emit more light (measured in lumens) per energy drawn (in watts) than inefficient light bulbs.

[26] Personal conversation with numerous data center operators, digital technology equipment suppliers, data center designers and other IT professionals between February 2001 and October 2001.

[27] Personal conversation with attendees at the "Austin Energy Model Data Center Energy/Design Meeting," February 12-13, 2001, Austin, TX; and personal conversation with attendees of Web Hosting Magazine. "Web Hosting Expo" Marriott Wardman Park, Washington DC. (August 20-22, 2001).

[28] Personal conversation with Jon Koomey, Staff Scientist, Lawrence Berkeley National Laboratory, October 12, 2001.

[29] John Cook. *Seattle Post-Intelligencer*. "Internet data gain is a major power drain on local utilities: Server farms' voracious appetite for electricity sparks several concerns" (September 5, 2000); and Thomas J. Grahame and David Kathan. *Broadband Wireless Business Magazine*. "Internet Data Centers & Electricity Growth: Will the Emerging Power Crisis Burn Your Bottom Line?" (July 17, 2001), Accessed October 18, 2001 at <http://www.shorecliffcommunications.com /magazine/volume.asp?Vol=12&story=147>

[30] Roger Anderson. *New York Times*. "Wattage Where It's Needed" (June 6, 2001); and Scott Harris. *The Industry Standard Magazine*. "Powering Up: How do you start an energy-hogging business in electricity-starved California? Bring your own power plant." (April 23, 2001). Accessed at <http://www.thestandard.com>

[31] Op. Cit. Jennifer D. Mitchell-Jackson

[32] Ibid.

[33] Ibid.

[34] Personal communication with Keith Reid, PG&E Senior Corporate Account Manager, November 8, 2001.

[35] Personal communication with Roger Duncan, Vice President Conservation, Renewables, and Environmental Policy. Austin Energy. October 5, 2001.

[36] American Gas Cooling Center. *Cool Times Energy Solutions*. " Natural Gas Technologies Promise to Feed Power-Hungry Data Centers" (Spring 2001); E Source. "Delivering Energy Services to Internet Hotels and Other High-Density Electronic Loads" Prospectus, E Source Multi-Client Study (Boulder, CO: 2001); Stuart M. Lewis. *Transmission & Distribution World*. "Utilities Cannot Afford to Become 'Sometimes Power & Light'" (April 2001); Jonathan Angel. *Network Magazine*.

"Emerging Technology: Energy Consumption and the New Economy" (January 5, 2001).

[37] Rich Miller. Website Editor, CarrierHotels.com. "How Much Power? Survey Shows Split Among Facility Operators." (March 21, 2001). Accessed August 15, 2001 at <http://www.carrierhotels.com/wiredspace /wired0321.shtml>

[38] Op. Cit. Jennifer D. Mitchell-Jackson; Op. Cit. American Gas Cooling Center.

[39] According to U.S. Energy Information Administration data, a typical U.S. home requires 1,370 W of electricity on average, thereby consuming approximately 12,000 kWh per year. <http://www.eia.doe.gov>

[40] See Table 4 of this paper for data references.

[41] Rich Miller. CarrierHotels.com "More Developers Adding Power Plants: Calif. crisis, utilities' demands for large deposits driving trend" (New York: Feb. 13, 2001). Accessed October 5 at <http://www.carrierhotels.com/news/February2001/imnpower0215.shtml>

[42] Op. Cit. Tadpole.com

[43] Op. Cit. Jennifer D. Mitchell-Jackson

[44] Op. Cit. Stuart M. Lewis

[45] Personal conversation with Keith Reid, PG&E Senior Corporate Account Manager, July 11, 2001; and Op. Cit. Jennifer D. Mitchell-Jackson

[46] Op. Cit. Jennifer D. Mitchell-Jackson

[47] Op. Cit. Jonathan Angel

[48] Here we use "balance of system" components to refer to all components of a data center other that the servers/computers and HVAC and chilling systems.

[49] Op. Cit. Jonathan Angel

[50] Personal conversation with Chris Wilkins, P.E., Hallam Associates, August 1, 2001.

[51] Op. Cit. Jennifer D. Mitchell-Jackson

[52] Ibid.

[53] Ibid.

[54] Op. Cit. Stuart M. Lewis

[55] Op. Cit. David Kathan and Thomas J. Grahame. (June/July 2001).

[56] Op. Cit. David Kathan and Thomas J. Grahame. (June/July 2001). In the case of PG&E, it can take up to 30 years to recover the capital investments they typically make for their system. Op. Cit. Keith Reid.

[57] Op. Cit. Rich Miller

[58] John Holusha. E-Business Alters ABC's of Real Estate. September 3, 2000. *New York Times*; And Jonathan Angel. Emerging Technology: Energy Consumption and The New Economy. January 5, 2001. *Network Magazine.*

[59] Personal conversation with numerous data center operators, digital technology equipment suppliers, data center designers and other IT professionals between February 2001 and October 2001.

[60] One study estimated electricity costs at 2% of data center income. Because energy is a very small cost of data center operation, and because billing departments rather than facility operators pay the energy bill, incentives to save energy are very low. Op. Cit. Jennifer D. Mitchell-Jackson

[61] Joe Romm. Center for Energy and Climate Solutions. "Cool Companies" (Washington DC: 1999).

[62] Active Hardware. "1.8Ghz Intel Pentium 4 Processor" (July 12, 2001). Accessed November 26, 2001 at <http://www.active-hardware.com/english /reviews/processor/p4-18.htm>

[63] Op. Cit. Jonathan Angel

[64] IBM. "IBM Launches "High-Powered" Initiative for 'Low-Power' Products and Services" (Armonk, NY; October 1, 2001) Accessed October 5, 2001 at <http://www.ibm.com/Press/prnews.nsf/jan/5328A987B028E83085256AD8004A3972>

[65] Transmeta Corporation. "Transmeta Recognized By U.S. Environmental Protection Agency For Energy-Efficient Crusoe Microprocessor" (Santa Clara, CA and Washington, D.C.; June 7, 2001) Accessed October 5, 2001 at <http://www.transmetazone.com/releaseview.cfm?releaseID=610>

[66] Transmeta Corporation. "Fibercycle Announces First Ultra-Dense Servers Using Transmeta Processor" (Los Gatos, CA: March 7, 2001) Accessed October 5, 2001 at <http://www.transmetazone.com/releaseview.cfm?releaseID=527>

[67] Op. Cit Tadpole.Com

[68] Russ Zabel. South County Journal. "`Server Farm' Boom Poses Threat to NW Power Supply" (Seattle: December 26, 2000)

[69] Op. Cit. Jennifer D. Mitchell-Jackson; Op. Cit. Keith Reid

[70] Op. Cit. Jennifer D. Mitchell-Jackson

[71] Op. Cit. Keith Reid

[72] Op. Cit. Jennifer D. Mitchell-Jackson

[73] Personal conversation with Stephen Spinazzola, P.E., RTKL Associates, September 7, 2001.

[74] Op. Cit. Chris Wilkins

[75] Yazaki Energy Systems, Model WFC-10, Farmers Branch, TX. Promotional material, 2001.

[76] EPRI. Consortium for Electric Infrastructure to Support a Digital Society (CEIDS) "Draft R&D Plan" (March 7, 2001)

[77] Dan Baer. Liebert Corporation. "Emerging Cooling Requirements & Systems in Colocation Spaces" Presentation at Colocation Summit March 14-16, San Francisco. (March 16, 2001)

[78] Richard Bourne. *Architecture*. "Night Moves" (November 1999). Accessed October 5, 2001 at <http://www.architecturemag.com/nov99/tp/green /green.asp>

[79] Robert F. Sullivan. Uptime Institute white paper. "Alternating Cold and Hot Aisles Provides More Reliable Cooling for Server Farms" (Uptime: 2000).

[80] W. Pitt Turner IV and Edward C. Koplin. Uptime Institute white paper. "Changing Cooling Requirements Leave Many Data Centers at Risk" (Uptime: 2000)

[81] Ibid.

[82] Uptime Institute white paper. "Heat Density trends in Data Processing, Computer Systems, and Telecommunications Equipment. (Uptime, 2000). Available at <www.uptimeinstitute.org>

[83] RTKL. Press release. " RTKL Partners With APW-Wright Line To Unveil 'The Tower Of Cool': Patent Pending Thermal

Management System May Revolutionize Design for High-Density Data Centers, Reducing Cooling Costs 15-20%" (Atlanta: September 10, 2001) Accessed September 10, 2001 at <http://www.rtkl.com/about/news.asp>

[84] Sanmina Enclosure Systems. Press release. "New Sanmina Product Dramatically Reduces Energy Costs in Data Centers" (Toronto: September 10, 2001) Accessed October 5, 2001 at <http://www.sanmina.com/company_info/press_releases/releases/pr09_10a_01.html>

[85] Op. Cit. Jennifer D. Mitchell-Jackson

[86] Op. Cit. Jonathan Angel. For more information see Northwestern's PACT website at <http://www.ece.northwestern.edu/cpdc/PACT/PACT.html>

[87] NASA Ames Research Center. "Technology Opportunity: Photonic Switching Devices Using Light Bullets" Accessed October 9, 2001 at <http://ctoserver.arc.nasa.gov/techopps/switching.html>

[88] Ibid.

[89] Lawrence Berkeley National Laboratory website, "Information Technology and Resource Use", accessed July 3, 2001 at <http://enduse.lbl.gov/Projects/InfoTech.html>; and Rachel Konrad. Server Farms on Hot Seat Amid Power Woes. CNET News.com. May 14, 2001. Accessed August 15, 2001 at <http://news.cnet.com/news/0-1007-200-5919632.html>

[90] According to the U.S. Energy Information Agency (EIA), total U.S. electricity consumption in 2000 was approximately 3,400,000 GWh. EIA Monthly Energy Review, "Electricity". Accessed October 17, 2001 at <http://www.eia.doe.gov/emeu/mer/elect.html>

[91] Op. Cit. Thomas J. Grahame and David Kathan. (May 2001).

[92] Op. Cit. Scott Harris; and Noam Levy. Mercury News. "Energy crisis raises new questions about massive Internet server farm" (San Jose: March 24, 2001)

[93] Op. Cit. Chris Wilkins

[94] Op. Cit. Scott Harris; and Op. Cit. Noam Levy

[95] Dan Whipple. Interactive Week. "Web-Hoster Uses Solar-Powered Server" (March 20, 2001) Available at <http://www.zdnet.com/intweek>

[96] Personal conversation with Steve May, President of SolarHost, October 5, 2001. For more information see <http://www.solarhost.com>

[97] Op. Cit. Jennifer D. Mitchell-Jackson

[98] Austin Energy. "Council approves chiller plant for Lamar-Sixth Austin MarketPlace" (Austin: February 17, 2000) <http://www.austinenergy.com>

[99] For more information on the Climate Savers Program see <http://www.cool-companies.org> and <http://www.worldwildlife.org/climate/climate.cfm>

[100] Center for Resource Management. Press release. "North American Telecom Industry Launches Environmental Charter" (Denver, CO: February 22, 2001). Accessed October 22, 2001 at <http://www.crm.org/crm%20html%20pages/telecom.html>

[101] For more information visit the Savings by Design website at <http://www.savingsbydesign.com>

[102] For more information on ESCOs visit the Energy Services Coalition website at <http://www.escperform.org> and the Energy Service Company Consortium website at <http://www.jri.co.jp/emergence/ESCO_e.html>

[103] U.S. DOE Office of Building Technology. "High Performance Commercial Buildings: A Technology Roadmap" (October, 2000) Accessed October 18, 2001 at <http://www.eren.doe.gov/buildings commercial_roadmap/>

[104] Edison Electrical Institute. "Power Companies, Developers Hold Second Conference on Internet Data Centers" *EEI Advisory.* (Washington: September 26, 2001); and Op. Cit. Stewart Lewis. Accessed October 16, 2001 at <http://www.eei.org/esg/na/nn/energy.html>

[105] Edison Electric Institute. "Energy News: EEI Launches 'Internet Hotel' Data Survey" *EEI National Accounts Network News*. (Issue #44, June 2001). Accessed October 16, 2001 at <http://www.eei.org/esg/na/nn/energy.html>

[106] For more information on the Consortium for Electric Infrastructure to Support a Digital Society visit their website at <http://ceids.epri.com/ceids/home1.jsp>

[107] For more information see the Uptime Institute white paper collection and industry specification web pages at <http://www.upsite.com/TUIpages/tuiwhite.html> and < http://www.upsite.com/TUIpages/tuifault_specs.html>

[108] The PCI Industrial Computer Manufacturers' Group develops hardware specifications for components that connect to a PCI (Peripheral Component Interconnect) bus. For more information visit their website at <http://www.picmg.org>

[109] For more information on the 7 x 24 Exchange visit their website at <http://www.7x24exchange.org/index.html>

[110] For more information on EPA's ENERGY STAR program visit their website at <http://www.energystar.gov>

[111] Rich Brown, Carrie Webber, and Jon Koomey. Lawrence Berkeley National Laboratory. "Status and Future Directions of the ENERGY STAR Program" Published in the Proceedings of the 2000 ACEEE Summer Study. (ACEEE: 2000) Also available as LBNL-45952 at <http://enduse.lbl.gov/EStar.html>

[112] The partnership between Verizon and EPA will be officially announced on or about November 16, 2001. Personal conversation with Betty Anderson, Verizon Energy Team. (October 17, 2001); Cool Companies. "Feature Story: Verizon--Dial "E" For Energy Savings" Accessed October 16, 2001 at <http://www.cool-companies.org/ads/featurestory.cfm>

[113] Personal conversation with Andrew Fanara, U.S. EPA ENERGY STAR Computers Program. (October 16, 2001).

[114] Op. Cit. Edison Electric Institute. (Issue #44, June 2001).

[115] Personal correspondence with William Tschudi, Lawrence Berkeley National Laboratory. November 15, 2001.

[116] Personal conversation with Martha Johnson, CBECS Manager, EIA. October 19, 2001. For more information see the CBECS website at <http://www.eia.doe.gov/emeu/cbecs/>

[117] EPRI "CEIDS R&D Plan Draft" (September 2001) Accessed at <http://ceids.epri.com/ceids/Docs/CEIDS_RD_Plan_092101.pdf>

The Renewable Energy Policy Project (REPP) supports the advancement of renewable energy technology through policy research. We seek to define growth strategies for renewables that respond to competitive energy markets and environmental needs. Since its inception in 1995, REPP has investigated the relationship among policy, markets and public demand in accelerating the deployment of renewable energy technologies, which include biomass, hydropower, geothermal, photovoltaic, solar thermal, wind and renewable hydrogen. The organization offers a platform from which experts in the field can examine issues of medium- to long-term importance to policy makers, green energy entrepreneurs, and environmental advocates.

REPP receives generous support from the U.S. Department of Energy, the Energy Foundation, the Joyce Mertz-Gilmore Foundation, and the U.S. Environmental Protection Agency, the State of New Mexico, the National Renewable Energy Laboratory (NREL), the Bancker-Williams Foundation, and the Oak Foundation.

To order REPP publications, contact REPP at (202) 293-2898.

REPP publications are available on the Internet at http://www.repp.org

© Copyright 2001 Renewable Energy Policy Project

Non-Profit
U.S. Postage Paid
Washington, DC
Permit No. 1398

REPP
Renewable Energy Policy Project
1612 K St., NW, Suite 202
Washington, DC 20006