ENERGY NEEDS IN AN INTERNET ECONOMY: A CLOSER LOOK AT DATA CENTERS

by Jennifer D. Mitchell-Jackson
May 2001

submitted in partial satisfaction of the requirement for the degree
Masters of Science in the Energy and Resources Group
of the University of California, Berkeley

Approved:

Daniel M. Kammen   Date
Associate Professor, University of California, Berkeley

Jonathan G. Koomey   Date
Staff Scientist, End-use Forecasting Group
Lawrence Berkeley National Laboratory

Received:

Kate Blake   Date
Graduate Assistant, Energy and Resources Group

Energy Needs in an Internet Economy:

A Closer Look at Data Centers

Jennifer Mitchell-Jackson

Energy and Resources Group, UC Berkeley

I. Introduction

Founded on the dispersion of the personal computer, the expansion of the nation's fiber optic network, and the spread of the World Wide Web, a "new economy" has emerged. But new technologies and new ways of conducting business are not the only changes. This Internet economy also brings with it new requirements for power and power systems.

In many ways, the Internet has facilitated a move toward more decentralized systems. Customers can now use the Internet to purchase products from stores around the country rather than being limited to stores in their neighborhood. Despite the decentralized nature of this web, the communication system involved with Internet transactions has led to a concentration of electricity demands at certain locations along the nation's fiber optic backbones. (See Figure 1.) At these locations, utilities face large electricity demands from facilities called data centers.

Figure 1. U.S. Data Center Hosting Facility Locations

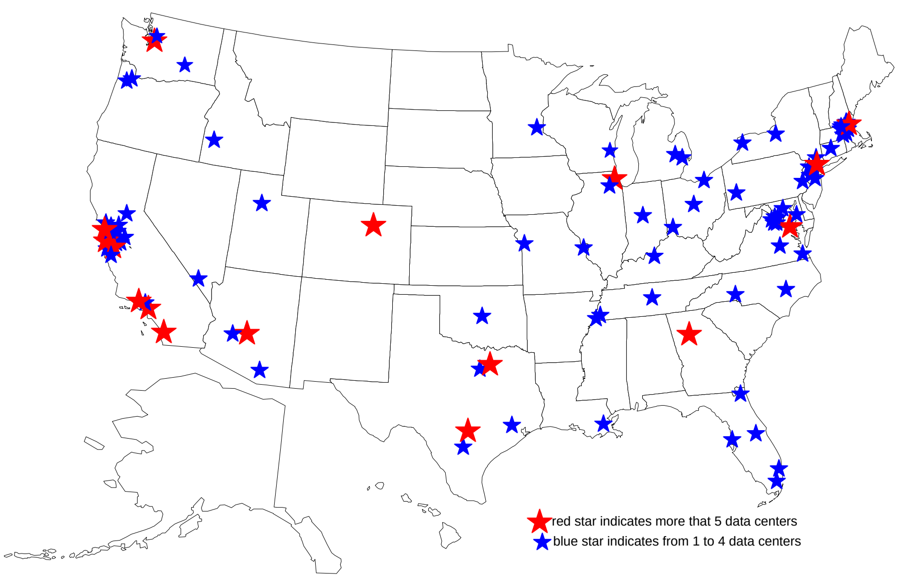

Source: Salomon Smith Barney

Data centers, also commonly referred to as server farms or dot-com hotels, are buildings that house computer equipment to support information and communication systems. The concentration of densely packed computer equipment leads to power demands that are much higher than those of a typical residence or commercial office building. Exactly how much power data centers require, however, is unknown. Although there are many estimates of the amount of electricity consumed by data centers, there are no publicly available *measurements* of power usage. Current assumptions about the power needs of these facilities are based on criteria that incorporate oversized, redundant systems, and several built-in safety factors. Furthermore, the estimates commonly cited in the media or in discussions with utilities assume that data centers are filled to capacity. These estimates, therefore, are greatly overstated.

In addition to requests for large amounts of power, data centers also require more reliable power than typically provided by the nation's electricity grids. While an electricity grid generally provides power with approximately 99.9% reliability, or around 9 hours of outages a year, data centers require what is being called "six nines," or 99.9999% reliability, which corresponds to approximately 32 seconds of outages a year. Moreover, data centers are demanding these requirements in a much shorter-than-usual time frame. While growth of this industry has recently slowed, these facilities had been popping up in 8 to 12 months in an attempt to minimize "time to market."

New requests for power have altered previous energy forecasts. One of the questions that this project seeks to answer is "Are these forecasts accurate?" Understanding the accuracy of these estimates is important because energy planning in the United States has been highly dependent on forecasted demand. Faced with new requests for power, utilities must decide how to respond. Building new infrastructure and acquiring new power resources to meet demand will have serious costs. If priced correctly, these costs can be recouped through electricity charges, but if demand is significantly overstated, the utilities will spend more on infrastructure than they could ever hope to recover. Furthermore, utilities may build infrastructure to meet the power demands of an industry that could fade as rapidly as it appeared.

The goal of my project is to provide a better understanding of the electricity demands of this industry and to help establish realistic growth projections for the future. My research project examines the energy consumption of data centers, delves into the industry standards and lack of consistent terminology that underlie the current projections, and attempts to add actual measurements and analysis of real-world data to the debate over how much energy these facilities require.

In conjunction with the Lawrence Berkeley National Laboratory (LBNL), I recorded energy indicators and electricity demands at one Bay Area facility, and examined electricity bills for several other facilities throughout the country. For confidentiality reasons, these data centers will be identified only by their approximate location and not by the companies that own them. In this report, I give some background on data centers, present the findings of my analysis, outline some utility-level policies that have been created to attempt to tackle this issue, and offer options for the future.

II. Data Center Fundamentals

Data centers provide the physical infrastructure for housing computer equipment, often referred to as information technology (IT) equipment.[1] A data center's main function is to provide guaranteed reliable power, security, cooling and connectivity to the rest of the Internet via a high-capacity backbone.

There are several types of data centers, each of which has its own unique characteristics. The main categories of data centers are: corporate data centers, co-location data centers, and managed data centers.[2] In the late 1980s and early 1990s, most data centers were corporate data centers. One example of a corporate data center would be a large data center owned by a bank for the purpose of storing financial records and providing software applications specifically for that bank. Over time, many of these corporate data centers began to be outsourced to "hosting facilities." Outsourcing a corporate data center provided the company with a form of insurance. If a fire, a natural disaster, or a power outage were to render a corporate office inoperative, hosting facilities offered alternative locations that could handle the company's critical functions.

Co-location and managed data centers are two types of hosting facilities. Exodus, HostPro, Globix, AT&T and UUNet are examples of companies that own and operate large hosting facilities. In co-location data centers, or "colos," an outside company such as Yahoo could rent space and bring in its own computer equipment. In these data centers, the customer-managed space is usually protected by wire cages that extend from floor to ceiling for security purposes. One data center computer room might contain many cages, each set up for an individual client. In "bare colos," the cages might be the only

[1] Data centers may also house conventional telecommunications equipment such as telephone switches but this investigation focuses only on the growth of Internet data centers. A detailed investigation of conventional telecommunication data centers is beyond the scope of this report.
[2] There are also university or government data centers but I am classifying those under the category of "corporate" or individual data centers.

structure within the room. In other co-location data centers, the facility might provide both cages and the racks to hold the computer equipment.[3]

In managed data centers, the owner of the data center owns not only the racks but also the computer equipment within the facility. These data centers also provide other services such as software management. From an energy standpoint, the important distinction is that the floor plan and layout for a managed hosting facility can be planned ahead of time while the equipment and actual power demands of a co-location facility are up to the customers that rent the space. The division, of course, is not as clear-cut as it may seem. The exact power needs of a managed hosting facility are not always known in advance since the rate of technological change in this industry is rapid. However, there are more unknown variables in a co-location facility.

There may also be physical differences between managed and co-location facilities. Managed hosting facilities are more likely to have additional staff such as computer technicians and software engineers to monitor and maintain the computer equipment. Since managed hosting facilities have more employees, these facilities will most likely have more office space.[4] In addition, since data center customers do not access the computers in a managed hosting facility, cages are usually unnecessary. Finally, a managed hosting facility can position computer equipment in the optimal configuration and ensure that no circuits are overloaded, whereas in a co-location facility, it is possible that some customers' cages may be filled with computer equipment, while other space is not yet rented. Co-location facilities, therefore, are more at risk of electrical imbalances or overloaded circuits.

Many hosting facilities have a combination of managed hosting and co-location space. Current estimates by Salomon Smith Barney indicate that there is roughly twice as much

[3] In addition to the physical infrastructure, some colos might also provide hardware and software management services. These are 'managed colos' and fall in between the two classes of data centers that I describe.

[4] Note that office space has a lower power density than computer rooms. The power density of a typical office space ranges from 5 to 10 W/ft^2.

co-location space as managed hosting space, and that co-location will most likely continue to dominate.[5]

Data centers come in many sizes.[6] While some data centers occupy one or two floors within a multistory office building, others occupy entire buildings. For hosting facilities, larger data centers allow for a greater number of customers and thus economies of scale. Large data centers tend to range in size from 20,000 square feet to 250,000 square feet. There are, however, some data centers that are even larger. One Exodus data center being built near Seattle, for example, is reported to be approximately 576,000 square feet.[7] There is also a proposal by U.S. Dataport to build a 2.2 million square foot data center complex in San Jose, California.[8] This area, however, describes the entire footprint of the building (everything from the computer space to hallways and bathrooms). The area devoted to computer equipment is much smaller than the gross area described above.

There are at least three estimates of total U.S. data center floor space. All of these estimates are based on surveys of major data center hosting companies in which the companies were asked the location and size of their facilities.[9] One of these estimates, an estimate by Richard Juarez of Roberston Stephens, surveyed 60 companies.[10] Juarez estimated the total gross floor area. This estimate, therefore, is likely to include significant amounts of office space, hallways or other areas that may or may not be used for activities related to the central computer room. As a result, it may not be representative of the area needed to support Internet computer equipment.

---

[5] Mahedy, Stephen and Dan Cummins and Danny Joe. "Internet Data Centers: If Built…Will They Come," Salomon Smith Barney Report, 3 August 2000. (Salomon Smith Barney Report)

[6] From now on, use of the term "data center" will refer to data center hosting facilities unless otherwise stated. The term data center hosting facilities does not include corporate data centers.

[7] Cook, John. "Internet data gain is a major power drain on local utilities," *Seattle Post Intelligencer*, 5 Sept 2000.

[8] Lazarus, David, "Net Complex a Dilemma for San Jose," *San Francisco Chronicle*, 22 March 2001. This is a compound of 10 buildings. Square footage from http://www.ci.san-jose.ca.us/planning/sjplan/eir/USDataport/US-Dataport-Text.htm.

[9] These numbers do not include corporate data centers.

[10] 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.

Two other groups have also estimated total data center area. Their numbers include net floor area, or the central area devoted to computer equipment. Salomon Smith Barney surveyed approximately 40 companies. Based on these surveys, they anticipated that the there would be nearly 10 million square feet of data center computer rooms built or under construction in the United States at the end of 2000, and that this number would rise to approximately 17.9 million square feet by the end of 2001.[11] The Yankee Group also estimated total U.S. data center space devoted to computer equipment.[12] They predicted that there would be 9.3 million square feet of computer rooms at the end of 2000—approximately 6% less than in the Salomon Smith Barney Report.

Figure 2. Computer Room Floor Space In U.S. Data Centers

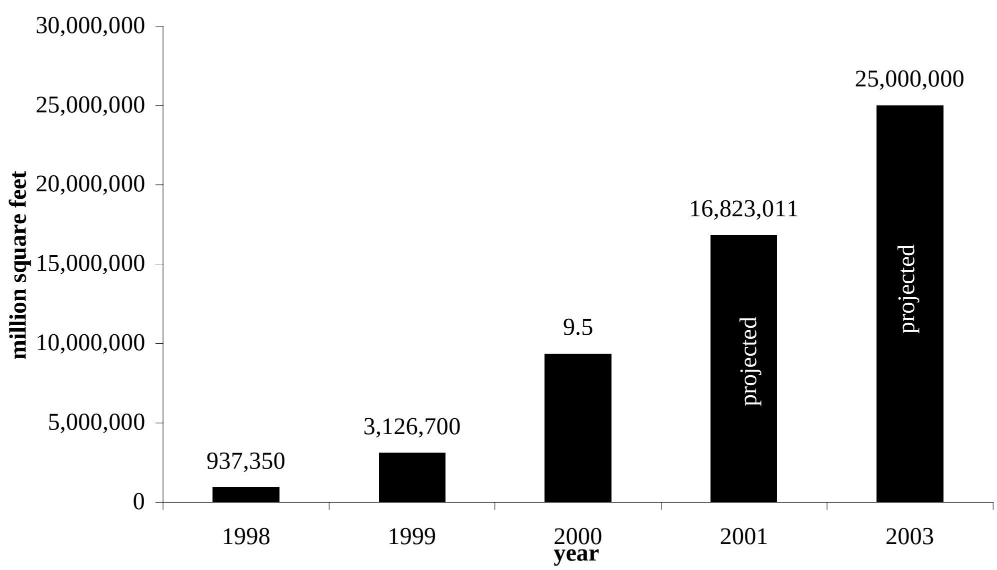

Source: Yankee Group, 2000; 2001 value estimated from the growth rate in the Salomon Smith Barney Report (Mahedy et al., 3 August 2000).

Based on the Yankee Groups' numbers for 1998 through 2000, and using the growth rate from the Salomon Smith Barney Report to make an estimate of the computer room area

[11] Their total estimates also included a significant number of international data centers but these are not included in the 10 million ft^2. Mahedy et al., 3 August 2000.

in 2001, Figure 2 shows that the net floor area of U.S. data centers has grown significantly. The total computer room area in U.S. data centers is estimated to jump from less than a million square feet in 1998 to approximately 17 million square feet by the end of 2001. Current projections indicate that it will continue to grow to over 25 million square feet by the end of 2003, although the recent economic events in this industry have led to a significant slowdown of data center construction.[13]

These numbers are rough estimates. The surveys did not include all of the companies that own data centers, thus some data centers have been excluded. It is also possible that these numbers are overstated. The Salomon Smith Barney report, for example, indicates that at the end of 2000, HostPro had five data centers with approximately 64,000 square feet of computer room space, as well as 94,000 square feet of data center computer room space in "unspecified" locations. According to the HostPro website, however, the company currently has only five data centers.[14] It is possible that HostPro intended to build additional facilities, but has not yet done so. It should be noted, however, that the HostPro estimate included in the Salomon Smith Barney survey might be more than twice as large as HostPro's actual computer room floor area.

According to the Robertson Stephens report, there were approximately 320 data center hosting facilities in the United States at the end of 2000.[15] As the map in Section I indicates, the majority of these hosting facilities are clustered in major metropolitan areas. There are 16 major cities around the country that have or have planned more than 5 data centers each. These locations are listed in Table 1 below. For co-location data centers in particular, it is important that the facilities are located near customers since the customers are in charge of placing the computer equipment in the data center and servicing or adding new pieces of equipment as needed.

---

[12] Yankee Group, "Executive Summary of The U.S. Collocation Market: High-Tech Real Estate Heats Up," 2000, www.yankeegroup.com/ viewed 28 March 2001.
[13] Projection from Lazarus, 22 March 2001.
[14] Information from the HostPro website, www.hostpro.com, viewed 12 April 2001.
[15] This was about 45 more than Salomon Smith Barney's projections but the Robertson Stephens report covered 60 companies while the Salomon Smith Barney report surveyed only 40.

Silicon Valley, California (consisting of four Bay Area counties: Alameda, San Francisco, San Mateo, and Santa Clara) is the largest single data center hub in the country. The Robertson Stephens report estimates that there are approximately 54 data center hosting facilities in this region.[16] According to this report, the hosting facilities in the Bay Area make up approximately 17% of all major hosting facilities in the United States. The Salomon Smith Barney report (which pre-dates the Robertson Stephens report and covers only 40 companies) agrees that approximately 15% of the data centers are being built in this region, but indicates that these data centers may represent less than 15% of the country's data center computer room floor area. The Salomon Smith Barney report also indicates that computer room floor space in the New York City area may rival the amount of floor space in the Bay Area.[17]

Table 1. Cities with more than five data center hosting facilities built or planned

Atlanta, GA
Austin, TX
Boston, MA
Chicago, IL
Dallas, TX
Denver, CO
Irvine, CA
Los Angeles, CA
New York City, NY
Phoenix, AZ
San Diego, CA
San Francisco, CA
San Jose, CA
Santa Clara, CA
Seattle, WA
Washington, DC/Northern Virginia

Of course, there is much uncertainty in the total number of data centers and the quantity of computer room floor area across the country. Companies in this industry come and go

[16] Juarez et al., 29 January 2001.
[17] Mahedy et al., 3 August 2000.

frequently. Changes due to mergers or the recent economic situation could significantly alter these numbers. Neither of these reports, for example, includes the proposal for the new U.S. Dataport complex described earlier. The siting of this facility, however, is in dispute. In the end, this data center complex, like other data centers, may or may not be built.

Currently, the best estimates indicate that there were approximately 9.5 million square feet of data center computer rooms in the United States at the end of 2000, and that this will reach approximately 17 million square feet by the end of this year. Some of the electricity demands from these facilities are due to new growth, but some unknown portion is due to consolidating several smaller corporate data centers from office buildings throughout the country into larger hosting facilities.[18] In fact, one major market for some smaller data centers is convincing Internet start-ups or other businesses to move their Internet equipment out of their janitor's closet and into the safe and secure environment of a co-location hosting facility. In these situations, the previously dispersed computer equipment is often moved less than ten miles from the office where it was originally located.

Since energy planning in the United States has been highly dependent on forecasted demand, understanding whether requests for power from data centers represent actual new demand is extremely important. The local and national energy implications of the growing number of data centers are discussed in the following sections.

[18] Further investigation of corporate data centers is beyond the scope of this report.

III. The Current Debate

When electricity and the Internet are mentioned together, it is easy to focus on data centers since they represent a single concentrated load from computer equipment that is on 24 hours a day, 365 days a year.[19] These loads add up, but to how much? Do data centers pose a national threat? Are data centers responsible for the current electricity crisis in California? The overall impression from the media is clear. In the midst of the current California energy crisis, bylines exclaim: "Data centers pose serious threat to energy supply.[20]"

The belief that data centers might be one of the main reasons for Bay Area power shortages is exacerbated by the confusion over Internet electricity consumption. In a report for the Greening Earth Society, a non-profit backed by coal interests, Mark Mills and Peter Huber claimed that the Internet is responsible for 8% of all national electricity consumption, and that all office equipment uses 13%.[21] Furthermore, Mills and Huber project that, within the next two decades, office equipment will account for half of all U.S. electricity use. This report was summarized in a 1999 issue of *Forbes Magazine* under the title of "Dig more coal—the PCs are coming.[22]" Jon Koomey and his colleagues from the Lawrence Berkeley National Lab (LBNL) have refuted the Mills and Huber numbers, offering revised estimates of the Internet's electricity consumption. Koomey and his colleagues have estimated that all office, telecommunications, and network equipment in the United States use only 3% of the country's total electricity consumption—a factor of four reduction from Mills's estimate of 13%. Despite this LBNL report, the media and many other sources continue to cite Mills's inflated values.

---

[19] Basler and Hofman determined a flat load is representative of the energy demand of these facilities. Their paper states that, "It was established that the recorded power was constant in terms of time and does not depend on the data-flow quantity or network topology." Basler and Hofman, "Energieverbrauch von Netzwerkkomponenten (English Version)," Bundesamt für Energiewirtschaft Forschungsprogramm Elektrizität, 26 November 1997.

[20] Byline from Bors, Douglas, "Data Centers Pose Serious Threat to Energy Supply," *bizjournal.com*, 6 October 2000.

[21] This included all home and office personal computers as well as server farms and other internet-related equipment.

[22] Mills, Mark and Peter Huber, "Dig more coal—the PCs are coming," *Forbes Magazine*, 31 May 1999.

Recent reports of enormous power demands by data centers reach far beyond California's Silicon Valley. The *Seattle Post Intelligencer* reported that data centers requested 445 MW of power in a small area near Southcenter Mall outside of Seattle.[23] And on the East Coast, a *New York Times* article from July, 2000 reported that "One developer looking at [a] site in North Jersey for the home of a potential million square foot data and communications center asked Public Service for 100 MW of power…one third of what [the utility] give[s] the whole city of Newark.[24]"

While the media has lent credibility to these mammoth power demands, the companies that own the data centers are making the large forecasts. According to the *Sacramento*

Figure 3. New Power Requested By Data Centers in PG&E Region

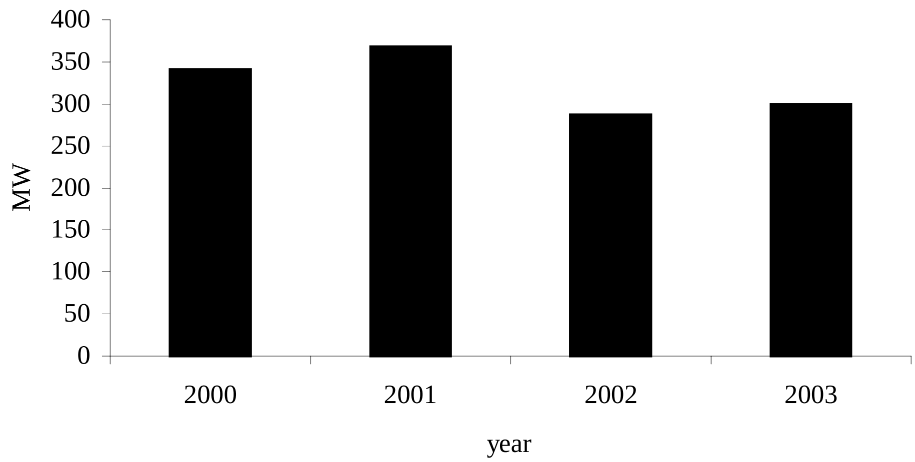

Source: Energy Solutions and Supersymmetry, Data Center Market Research Study, Presentation for PG&E Customer Energy Management, 19 October 2000.

Bee, 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

[23] Cook, 5 September 2000.
[24] Feeder, Barnaby, "Digital Economy's Demand for Steady Power Strains Utilities," *New York Times*, 2 July 2000.

average year.[25] And Keith Reed, a Senior Corporate Account Manager with PG&E, has indicated that data center customers in PG&E's territory are forecasting unheard of leaps in electrical demand.[26] PG&E has 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.[27] (See Figure 3.)

Since the largest data center hub lies in the midst of the California power crisis, it is understandable why people might implicate data centers as a threat to power security. But while there is clearly a stated demand from data centers (see Figure 3), actual demands have not yet materialized. Despite claims by a recent *Computer World* article that power demands "skyrocketed by 12%" in the heart of Silicon Valley compared to a

Table 2. Electricity Consumption in Silicon Valley Versus California

Year	**Silicon Valley Total Electricity Consumption**		**Statewide Total Electricity Consumption**	
	in million kWh	% growth	in million kWh	% growth
1990	31,436		228,038	
1991	31,140	-1%	222,260	-3%
1992	31,587	1%	226,988	2%
1993	31,679	0%	227,624	0%
1994	31,467	-1%	230,097	1%
1995	32,220	2%	230,990	0%
1996	32,911	2%	239,168	4%
1997	34,469	5%	246,225	3%
1998	34,289	-1%	249,926	2%
1999	35,360	3%	252,267	1%
	Overall Growth 1990-1999	12%	**Overall Growth 1990-1999**	11%
	Average Annual Growth	1.3%	**Average Annual Growth**	1.1%

Source: California Energy Commission website, "Silicon Valley Electricity Consumption," http://38.144.192.166/electricity/silicon_valley_consumption.html, viewed 29 March 2001.

[25] Peyton, Carrie, "Data servers crave power: High-tech electricity needs amplify crisis," *The Sacramento Bee*, 26 November 2000.
[26] Peyton, 26 November 2000.
[27] Energy Solutions and Supersymmetry, Data Center Market Research Study, Presentation for PG&E Customer Energy Management, 19 October 2000.

statewide growth rate of 2% or 3%, the California Energy Commission's webpage on "Silicon Valley Electricity Consumption" (seen in Table 2) shows that total electricity use in the Silicon Valley has not grown at a higher rate than in the rest of California.[28] Annual electricity growth in both areas has averaged just slightly over 1%. State electricity growth, in fact, has just kept pace with population growth, which has also grown at approximately 1% per year.[29] A comparison of non-residential growth in these two areas shows approximately the same trend over this period: Silicon Valley non-residential consumption grew by 11% and statewide non-residential consumption grew by 10%.

Overstated power demands have both helped and harmed data centers. On one hand, data centers exaggerate these numbers as a marketing tool, in effect claiming that they are bigger and better than their competitors. On the other hand, recent media attention has cast data centers as villains in the energy crisis. Clearly, there is a debate over how much energy these facilities use. In the search for answers, the first barrier that must be overcome is the lack of standard terminology and consistent metrics.

[28] Hall, Mark, "Net Blamed as Crisis Roils California," *Computer World,* 15 January 2001. California Energy Commission website, "Silicon Valley Electricity Consumption," www.energy.ca.gov/silicon_valley_consumption.html, viewed 29 March 2001.

[29] Brown, Richard and Jonathan Koomey, "Analysis of California Electricity End-Use (draft)," Lawrence Berkeley National Laboratory, May 2001.

IV. Defining Common Metrics

Power in data centers is commonly discussed in terms of power density (in watts per square foot or W/ft^2). It is usually unclear, however, what this watts-per-square-foot number means. Often, which equipment is drawing the power (in watts) and the footprint of the area that is being referred to (in square feet) are not clearly stated. The lack of common metrics leads to confusion because individuals discussing the power demands of these facilities may be comparing two dissimilar values. Furthermore, confusion about power densities has led to inaccurate calculations of data center power needs. In an attempt to give the reader a better understanding of this issue, I describe the configuration of a typical data center in relation to 1) the footprint and 2) the power load.

The Footprint

As stated earlier, the size of a data center hosting facility can vary significantly. Some facilities occupy entire buildings, while others may occupy only one or two levels in a multistory office building. The total floor space for the entire facility is often referred to as the gross floor area.

Within data centers, part of the floor space is devoted to computer and Internet hosting equipment (e.g., servers, routers, switches, tape drives, etc.), while other space is partitioned for office space, fire suppression equipment or other auxiliary equipment. (See Figure 4.) In an average hosting facility, approximately 50% to 60% of the facility's footprint is designated for computer equipment.[30] This number, however, may be significantly lower if the facility is a multipurpose facility with some areas being used for office space or other commercial uses. For example, the Globix website states that Globix operates two data centers. One of these data centers is 61,000 ft^2 and the other is 160,000 ft^2. In both facilities, however, the area designated for computer equipment is 24,000 ft^2.[31]

[30] The average net to gross ratio found in the Salomon Smith Barney report was 0.6, or 60%.
[31] Globix, www.globix.com, viewed 15 March 2001.

Figure 4. Sample Layout of a Data Center

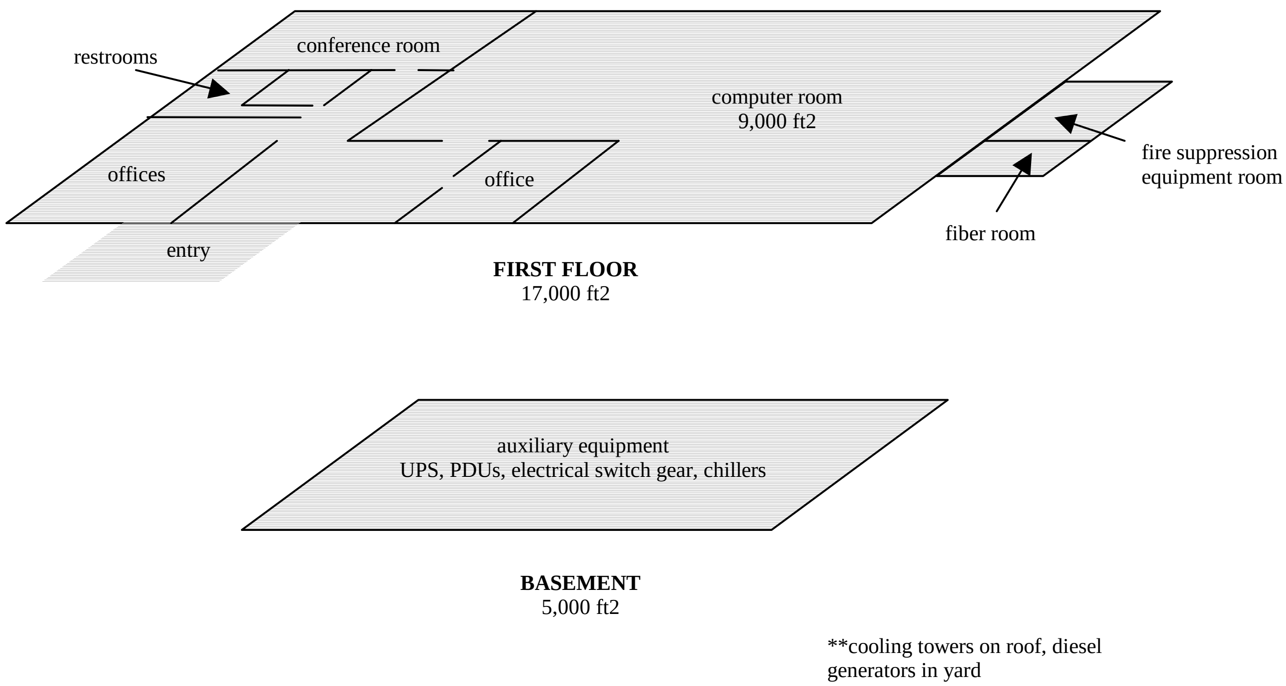

The main computer room area is commonly referred to as the net floor space. This area is also often called "raised floor area" due to the fact that the floor is usually elevated so that cool air can be blown in from underfloor air ducts. Within this area, 20% to 33% of the floor space is usually covered with rows of equipment racks or cabinets (i.e., enclosed racks). The low coverage of this area is due, in part, to building codes that require three-foot aisles between racks. Approximately 50% of the area in the computer room is inactive aisle space or service clearance area. The remaining 20% is usually for support equipment such as computer room air conditioning units (CRAC units) that must currently be located near the computers. Other mechanical or electrical equipment (mentioned in the power section below) is usually kept in external rooms in order to leave as much rentable space as possible.

The storage racks that hold the computer equipment are usually 6 to 8 feet tall with a footprint of approximately 6 square feet (2 ft x 3 ft).[32] (See Figure 5.) While the computer equipment that fills these racks can range in size, the trend is towards smaller computers that are approximately 1.75 inches high, about the size of a pizza box. The standard designation for the height of the computer equipment is "U," where 1U = 1.75 inches. A large percentage of the equipment, however, is still larger than this minimum size of 1U. The average desktop computer, for example, is 4U.

Figure 5. Data Center Racks

Photo by Bruce Nordman of LBNL, 21 December 2000.
The computers shown in this photo are approximately 4U.

On a smaller scale, a recent paper by the Uptime Institute referred to the power density of a single rack.[33] While this is less common, computer manufacturers or facility managers describing a single overloaded rack may state a W/ft^2 value for a single rack. This can be further confused by the fact that the majority of the computer equipment fits horizontally

[32] Exact dimensions are closer to 19" x 32". Most racks tend to be 6 feet tall since configuring equipment in taller racks would require ladders.

[33] The Uptime Institute, "Heat Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment (version 1.0)," White Paper issued by the Uptime Institute, 2000, available at http://www.uptimeinstitute.org/heatdensity.html.

into the racks (as opposed to the alternative tower configuration) so the footprint of a computer and the footprint of a rack are approximately equal.

When stating a W/ft^2 value, the area (in ft^2) may refer to: 1) a single piece of computer equipment, 2) a single rack, or 3) the footprint of all of the racks excluding aisle space, 4) the footprint of the central computer room including aisles between racks but excluding exterior mechanical rooms and office space, 5) a single data center floor within a multipurpose building or 6) the footprint of the entire building. (See Figure 6.) Clearly, the implications of an isolated rack that draws 100 W/ft^2 are much less than if the average power density over every foot of the entire building is 100 W/ft^2. Unfortunately, however, numbers such as 100 W/ft^2 are often cited; and it is usually unclear which area is being discussed.

Figure 6. Defining Common Metrics: Alternative Footprints (in square feet)

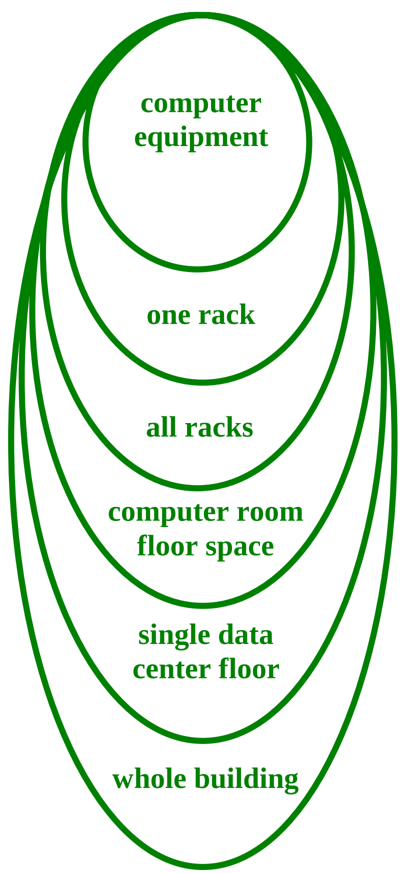

Power Loads

All currently operating data centers are connected to the utility grid. The loads from these facilities are assumed to be approximately steady because the computer equipment runs 24 hours a day, and weather is assumed to play a minor role compared to the heat generated by the computers.[34] Providing power in data centers is critical, and most data centers build all of their internal systems with redundancy. This is typically referred to as an N+1 design; if one piece of equipment fails, there is always a backup. If, for example, a facility needs two generators to support its power needs, the facility would have three. And some systems go beyond the N+1 design to a 2N design. In this case, the back-up system is also backed up. Many data centers also require this redundancy on the utility-side of the meter. They request duplicate transformers and substations from the local utility to ensure reliability. In addition, these facilities have on-site back-up diesel generators that will provide power to the data center in the case of an outage from the grid.

Power (in watts) can be described at several different levels. At the most general level, 480 volt (480V) power from the utility grid flows through a meter at the building envelope. From this meter, inflowing electricity is split between 1) the uninterruptible loads including the computer equipment, and 2) the other loads that, while important, could be disrupted temporarily without major consequence. See Figure 7 for a simplified schematic drawing.

The uninterruptible loads usually flow through several uninterruptible power supplies, or UPSs. A UPS smoothes dips in incoming voltages to provide a more uniform power source to the computers. It also acts as a battery in the event of an outage and provides power to the computer equipment while back-up generators are started. In the most reliable UPSs, all incoming AC power is converted to DC power to feed a battery, and the DC from the battery is converted back to AC to feed the racks of computer equipment. This continuous double conversion process reduces the chances that the

[34] See footnote 19.

Figure 7. Electricity Flows in Data Centers

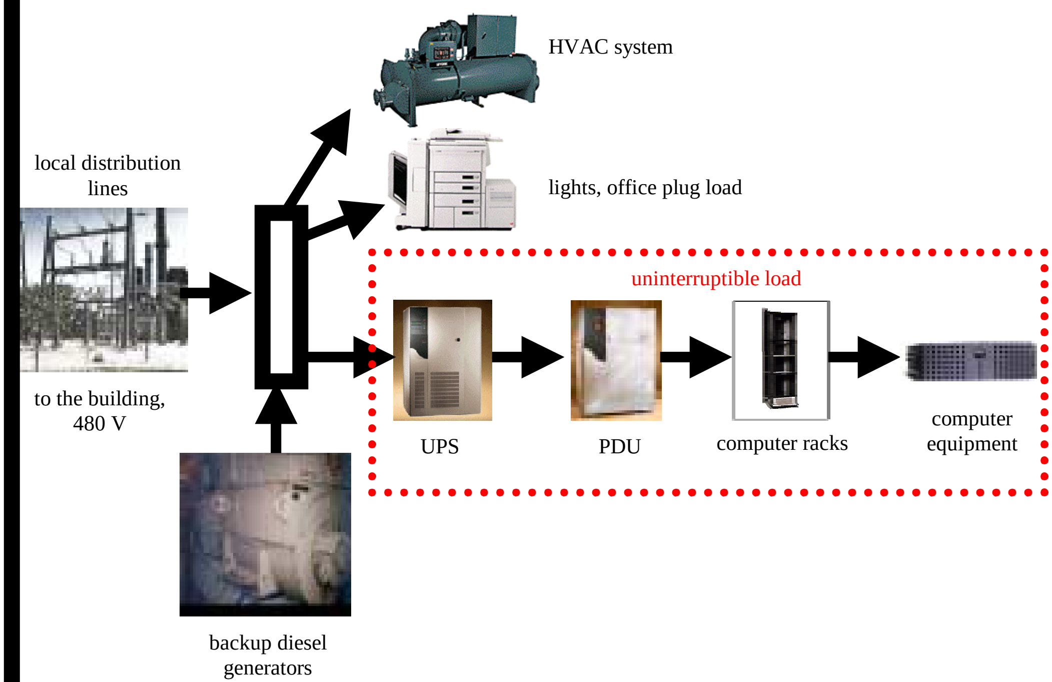

computer equipment will lose power. While it is possible to feed the racks of computer equipment directly from the AC grid while making sure that the battery is charged, in this type of system a power outage might cause the computer equipment to lose power for a fraction of a cycle while the power source switches from the grid to the batteries.

From the UPS, the electricity flows to several power distribution units (PDUs) that transform the electricity from 480V to 120V and send it to the appropriate panels.[35] Wires from these panels usually run several 120V/20A circuits to the racks where the computer equipment is plugged in. From the plug, the power is converted from AC power back to DC power in the computer. The DC power is consumed by the computer equipment in the racks and released as heat.[36]

[35] It may sometimes be converted to a different voltage such as 220V. This will depend on what the computer equipment requires.
[36] The Uptime Institute, 2000. It is interesting to note that the incoming electricity goes through a triple conversion process from AC to DC to AC and back to DC before being consumed.

All other loads, such as the heating, ventilation and cooling system (HVAC system), lighting, and other office equipment, are not routed through the UPS system. While these loads are usually on back-up generators, short outages do not lead to losses of critical information.

In summary, there are five general areas that require power: computer equipment, HVAC equipment, other auxiliary equipment to support the computer room (such as UPSs and PDUs), lights, and the office plug load (including copiers and faxes). (See Figure 8.)

Figure 8. Defining Common Metrics: Alternative Power Loads (in watts)

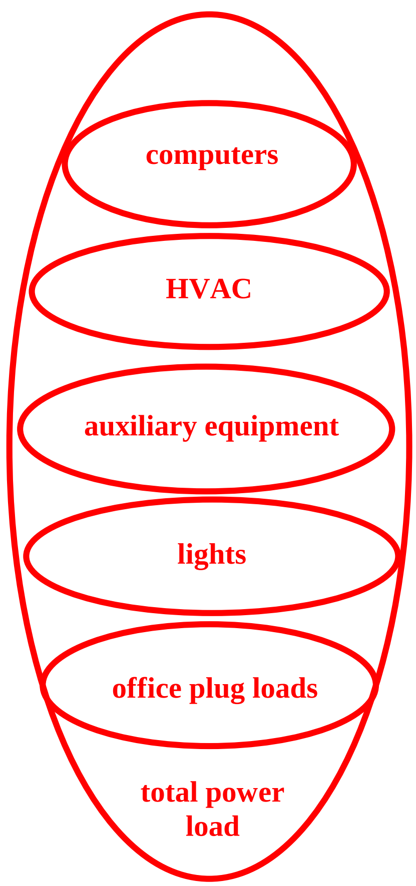

When stating the power density (in W/ft^2), designers and data center representatives (as well as the media) most commonly refer to the *computer power density*: the watts drawn by the computer equipment divided by the central computer room floor area. This value does not include HVAC, lighting, or other support systems. (See Figure 9.) Alternatively, utility representatives commonly refer to the *building power density*, or the total power load divided by the footprint of the whole building. The more office or low power density space the building includes, however, the lower the building power density. Thus, it is difficult to compare the power used by data centers if the buildings are different sizes or if the ratio of computer rooms to other space (such as office space) varies from building to building.

Figure 9. Defining Common Metrics: Alternative Terminology

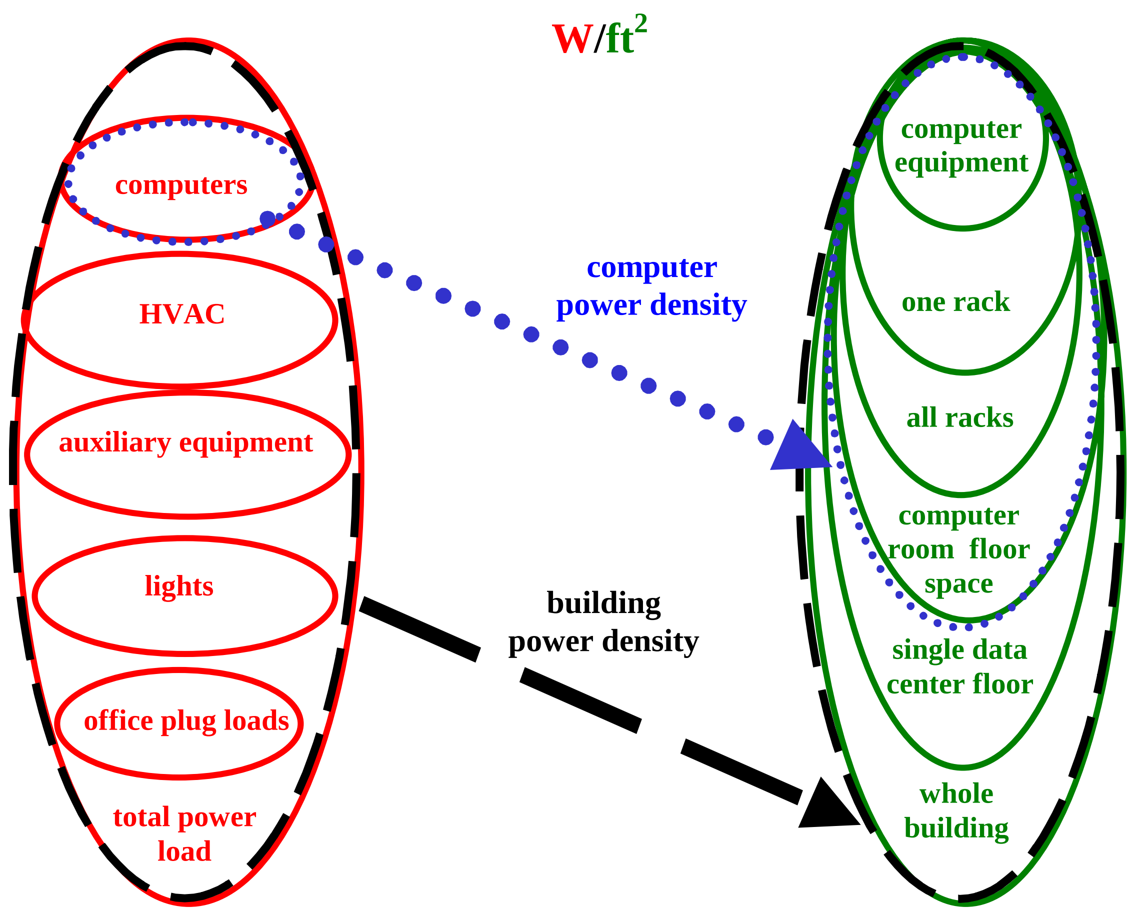

Data center designers indicate that the majority of existing data centers have a *design* computer power density of 50 to 75 W/ft^2.[37] Despite reports from facility managers that actual power densities are less than design, the companies building data centers have started to request much higher power densities and thus designers have begun to create data centers that can accommodate 100 to 150 W/ft^2 computer equipment loads in this central computer space.

The media and others discussing the energy use of data centers—usually unknowingly—estimate the power load of a data center facility by multiplying this design computer power density (e.g., 75 W/ft^2) times the floor area of the entire facility. This is not, however, an accurate estimate of power needs. As an example, in a 200,000 ft^2 facility with 50,000 ft^2 of core area designed for 75 W/ft^2, if the 200,000 ft^2 footprint of the building were multiplied by the 75 W/ft^2 design value, this would result in an estimate of 15 megawatts (MW). Not only does this calculation neglect to include the power used by lighting, HVAC and other systems, but it also assumes that the entire facility is full of IT or computer equipment when only one-fourth of the building (the 50,000 ft^2) is actually designated for this use.

The erroneous calculation performed above is one reason why the power needs of these facilities are overstated. Section VII discusses additional reasons for exaggerated power needs. In order to arrive at more accurate assumptions, however, it is critical that people explicitly state whether the power density that they are discussing is a design criteria or an actual measured value, whether it includes all power uses or just the computers, and whether it is for the entire building floor area or just the computer room.

Neither the two terms mentioned above—the computer power density nor the building power density—is really representative of the additional power loads of data centers, however. The computer power density neglects to include the additional power needed for HVAC, lights, and other support systems while the building power density may be significantly underestimated if the building includes lots of office space or other space

[37] Visit with a design team at Excite@Home, 20 November 2000.

not critical to the needs of the central computer area. A power density that includes all electricity used by the computer equipment as well as the electricity used by the necessary supporting equipment such as PDUs, UPSs, HVAC and lights is more representative of the true power needs of computer rooms. In this report, I refer to this as the *total computer room power density*. While the total computer room power density is a difficult number to grasp, it is the most representative of the requirements of this type of growth because it allows for comparisons between buildings of different sizes as well as between data centers at different stages of development. Table 3 below summarizes these three key terms.

Table 3. Key Terms For Discussing Power in Data Centers

Term	Definition
Computer Power Density	Power drawn by the computer equipment (in watts) divided by the computer room floor area (in square feet)
Total Computer Room Power Density	Power drawn by the computer equipment and all of the supporting equipment such as PDUs, UPSs, HVAC and lights (in watts) divided by the computer room floor area (in square feet)
Building Power Density	Total power drawn by the building (in watts) divided by the total floor area of the building (in square feet)

In the next section, I use these three terms to describe my estimates of the power consumed by a Bay Area data center in order to give a better sense of true electricity loads and how they are broken down within a data center hosting facility.

V. Estimating Data Center Loads

There are currently no publicly available measurements of electricity consumption in data centers. Because of security issues, it is difficult to gain access to data center hosting facilities, or to gather the relevant data required to estimate power densities and total loads. Data centers are often unwilling to share information because they feel it may compromise proprietary information. Through contacts at LBNL, I was able to gain access to a Silicon Valley data center and estimate the power needs of the systems and equipment within the building.

The estimates below are based on visits to the facility, measured data that I recorded, electrical and mechanical drawings for the facility, equipment counts and the manufacturer's specifications for the equipment found at this data center. I was also able to review electricity billing data for this facility. Further analysis of the billing data is discussed in the next section but information from the bills helped guide several estimates made below. While still rough, my estimates provide a benchmark for how electricity needs are broken down within data centers. This section also offers a sense of the complexities involved with estimating the power needs of these facilities, and gives some insight into areas where there may be opportunities for energy savings.

General Description of the Facility

The data below were collected from an approximately 125,000 square foot facility located in Silicon Valley. Like many data centers throughout the country, this data center was built within a renovated building in order to minimize construction time.

The basement of this facility was a large auxiliary equipment room containing batteries, UPS systems, PDUs, and switchgear for routing incoming electricity. A break down of the floor area is shown in Table 4 below. Additional equipment rooms with PDUs and fire suppression equipment were located on the first floor.

Table 4. Description of the Facility

Floor	*Total Area*	*Computer Rooms*	*Prior Use*	*Equipment Rooms*	*Office Space*	*Other Area*
(units)	*(ft²)*	*(ft²)*	*(ft²)*	*(ft²)*	*(ft²)*	*(ft²)*
Basement	27,165	0	0	25,710	0	1,455
First	43,400	27,500	0	4,240	0	11,660
Second	42,160	0	0	2,200	14,300	25,660
Third	12,600	0	12,600	0	0	0
Total	125,325	27,500	12,600	32,150	14,330	38,745

❖ Note that ‘Other Area’ includes computer rooms that are currently under construction.

At the time that I recorded my measurements, the first floor of the facility contained three active computer rooms with a combination of co-location and managed hosting space. The active computer rooms on the first floor occupied 22% of the total building floor area. (See Figure 10.)

The second floor of the facility contained approximately 14,300 ft^2 of office space and some equipment rooms. The remaining second floor area was still under construction at the time of my investigation. Approximately 19,000 ft^2 on the second floor was designated as a future computer room but was not yet occupied. The remaining second floor space was designed to be future mechanical and electrical equipment rooms or electrically inactive areas. Completion of this second floor data center will occur in a later phase of the development. (Note that the computer rooms that were under construction are included in "Other Area." Bathrooms, hallways, and lobbies are also included in the "Other Area" category.)

The third floor of this building remained in its prior use. All equipment on the third floor was in this facility prior to the recent renovation. Thus, the power loads from this part of the building do not represent new power requirements due to growth of the Internet, or the addition of a data center. This "Prior Use Area" represented approximately 10% of the total facility floor space.

Figure 10. Break Down of Facility Floor Space

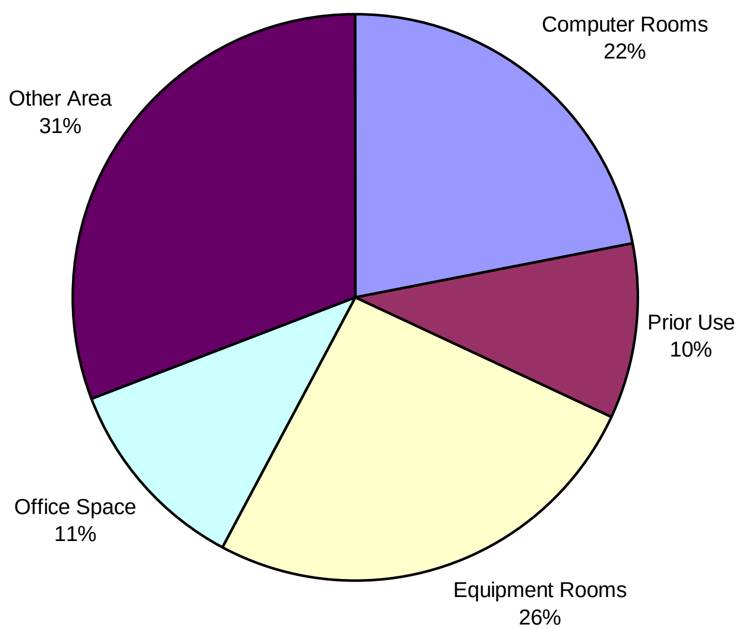

In the first floor computer rooms, all of the space was leased; however, on average, only one-third of the racks were full. I inventoried all of the equipment in an easily accessible portion of one computer room in order to get a sense of the type of computer equipment that was currently in this data center.[38] The area was part of a co-location facility where the computer equipment was enclosed in cabinets (racks with doors). These cabinets were located in an area that covered approximately 2,600 square feet. I selected this area because the equipment could be easily viewed and counted. Racks within cages were not selected because I did not have access to the cages and it was impossible to accurately count the equipment from outside the cages. The data, therefore, may be biased since the customers that rent cabinets, rather than cages, tend to be smaller customers. The racks in this area also tend to be more densely packed. The inventoried data for this area are reported in Table 5.

Table 5. Inventory of Equipment Found in Cabinets in a Co-location Hosting Facility

Type of Equipment	Number	Space in 'U's (where 1U=1.75")	Percent of utilized rack space devoted to equipment
Servers	229	596	61%
Switches	101	177	18%
Disks	18	88	9%
Routers	13	81	8%
Firewall	8	15	2%
Other	12	19	2%
Total	**381**	**976**	**100%**

- Data collected by Bruce Nordman and Jennifer Mitchell-Jackson, November 2000.
- The above equipment was located in approximately 2,600 ft^2 of a computer room.

Approximately 47% of the racks in this space were empty. The remaining racks had varying amounts of equipment. Servers, ranging in size from 1U to 8U accounted for 61% of the utilitized rack space. One third of these severs were 2U servers. While the data in Table 5 give a sense of the types of equipment in this space, I did not try to estimate power consumption based on this information because the energy demands vary depending on the internal configuration of the equipment. While servers generally use less power per unit area than routers, one 4U server may require significantly more power

[38] Bruce Nordman, LBNL, helped to identify the equipment.

than another 4U server if it contains more computer chips or is designed to perform more functions. It is difficult to determine the power requirements from the external appearance of the computer equipment.

1. Determining Power Demands From Computer Equipment

For each of the three computer rooms, eight power distribution units (PDUs) transformed the incoming 480V power to 120V.[39] Each PDU contained a system-monitoring panel that displayed three-phase volts and amps. I visited the Bay Area data center in November 2000, in January 2001, and in April 2001 to record the voltages and currents for the PDUs that served these computer rooms. Table 6 below is a summary table of my January measurements.[40] I arrived at the apparent power (in VA) by multiplying the average of the voltage (approximately 120V) by the sum of the currents (in amperes) from the three phases.

Table 6. Summary Table of PDU Data

Computer Room	Average Voltage	Total Current	Apparent Power Consumed By Computers	Real Power Consumed By Computers
	Volts	Amps	kVA	kW
One	120	868	104	101
Two	120	1590	190	184
Three	120	1257	151	146
		Total	445	432

- ❖ Data collected by Jennifer Mitchell-Jackson, January 2001.
- ❖ Stated amps and power values do not include the energy consumed by the PDU.
- ❖ Data for computer room one and three are inferred from readings of the 480V power. See Appendix A for more detailed data from the PDU data collection.
- ❖ A power factor of 0.97 was used to convert from apparent to real power. (See text for explanation.)

[39] In computer rooms one and three there were four active PDUs and four redundant PDUs. In the second computer room, most of the racks were connected to two PDUs to ensure redundancy, and all PDUs were in use. Detailed data from the PDUs is included in Appendix A.
[40] I chose to use the January data because billing data were also available for this month (see Section VI).

From these measurements, I determined that the computer equipment required approximately 445 kVA (apparent power). To convert apparent power to real power (in kW), I multiplied the apparent power by the power factor of 0.97 for the computer equipment. The power factor for newer computer equipment is usually high or close to unity because the computer equipment is corrected to eliminate harmonic distortions that might cause disruptions. New switching power supplies for Sun computers, for example, have active power factor correction to at least 0.99 in most cases.[41] Measurements from a both a New York City data center and an Oakland data center, however, indicated that the aggregate power factor for this equipment is closer to 0.97.[42] By multiplying the apparent power by the power factor, I determined that the computer equipment in these data centers required approximately 432 kW of power. Since the central computer room area was approximately 27,500 ft^2, the actual computer power density was slightly less than 16 W/ft^2.

2 Determining Power Used By The Prior-Use Area

From historic billing data, I was able to determine an approximate average power density for the "Prior-Use" area, which represented 10% of the total building area. The power density was approximately 20 W/ft^2 over this 12,600 ft^2 area. This value includes all of the equipment, lights, fans and plug loads for the third floor but does not include the power needed to provide chilled water to the air conditioning units (i.e., the central plant requirements.)

[41] Anonymous email from Sun Microsystems technical support desk, 07 February 2001. Email text as follows: an active power factor correction to at least 0.99 "has become a regulatory requirement for almost all new switching power supplies, with particularly strict regulations in the European Union. Uncorrected power factor can cause core saturation in the distribution transformers leading to early failure and decreased efficiency in the distribution grid." Also supported by The Uptime Institute, 2000.

[42] NYC data was collected by the facility manager at the data center from January 2000 through March 2000. The power factor readings were from the UPSs with loads. I collected the Oakland data on 21 December 2000 from the monitors on the active PDUs.

3. Determining Power Used By Computer Equipment In Office Space

The office space on the second floor of this facility also contained some computer equipment. There were fewer computers in this office space, however, than in an equally large commercial office space since the main employees of the building were mechanical and electrical staff. Some co-location customers also occasionally occupied the office space. During my visits, there were approximately twelve computers on in this space. The average heat gain for a typical office computer is approximately 55 watts.[43] A medium sized monitor would add an additional 90 watts.[44] This estimate is for an active computer and does not take into account that the computer and monitor would draw less if it is in a power saving mode, nor the fact that these computers are probably not on 24 hours a day. Furthermore, the laptops used by the co-location customers would require less power and release less heat than a desktop computer, but as an average estimate, I assume that a typical computer with a monitor uses 145 watts at all times. For twelve computers, this is approximately 1,740 watts, or 0.1 W/ft^2 over the 14,300 ft^2 office space.

4. Adding Lighting

In a typical commercial building, lighting accounts for about 1.8 W/ft^2 in office space.[45] I used this value to calculate the amount of power needed for lighting in the 14,300 ft^2 of office spaces on the second floor. Computer rooms and the remaining other space were less well lit than the office space. I did not have the electrical drawings for the lighting on the first floor, but from the second floor drawings, it appeared that the power density of the lighting in the computer rooms was approximately 1.1 W/ft^2. I used this value to calculate the power needs for the lighting in the computer rooms as well as the lighting in "Other Areas." This value is approximately the power density of lighting in the lobby

[43] Wilkins, Christopher and M.H. Hosni, "Heat Gain From Office Equipment," *ASHRAE Journal*, June 2000.

[44] Wilkins and Hosni, June 2000. Kawamoto, Kaoru and Jon Koomey, Bruce Nordman, Mary Ann Piette, and Richard E. Brown, "Electricity Used by Office Equipment and Network Equipment in the U.S.: Detailed Report and Appendices," LBNL Publication 45917, February 2001.

[45] Richman, Eric E. and Carol C. Jones, and JoAnne Lindsley, "An Empirical Data Based Method for Development of Lighting Energy Standards," Journal of the Illuminating Engineering Society, Summer 1999.

areas of a typical building.[46] Mechanical and equipment rooms tend to be slightly lower, thus I used a typical value of 0.7 W/ft^2 for these areas.[47] The total load from lighting was approximately 117 kW.

5. Adding In Other Loads

In addition to lights and computers, other office equipment such as copiers and faxes contribute small power loads throughout the office space. A recent *ASHRAE Journal* reported the heat gain to be approximately 1,100 watts from an office copier, 30 watts from a facsimile machine, 25 watts from an image scanner, and 550 watts from a large office laser printer.[48] These number, however, do not take into account power saving modes or end-of-the-work-day shutdowns. In a study that factored in power saving modes and shutdowns, Lawrence Berkeley National Laboratory found that a typical office copier draws an average of 100 watts and a large office laser printer draws an average of 32 watts over the course of a year.[49] These numbers give some reference points for calculating the additional loads in this space.

A study by Wilkins and McGaffin examined the office space in five buildings and measured the total heat gain to be between 0.44 W/ft^2 and 1.05 W/ft^2.[50] The office space examined in the Wilkins and McGaffin paper, however, was fully occupied and highly automated with a computer and monitor at every workstation. For my calculations, I assumed that this additional equipment drew just 0.3 W/ft^2 since the power density of the computers in this area was already included in an earlier step, and since this space was not densely occupied. In addition, 0.1 W/ft^2 was added to all "other" areas to account for small miscellaneous loads. These values carry with them less certainty than the measured data reported above, but they are small in comparison to the larger loads of the computers and HVAC system (discussed below).

[46] Richman et al., Summer 1999.
[47] Richman et al., Summer 1999.
[48] Wilkins and Hosni, June 2000.
[49] Kawamoto et al., February 2001. These numbers are averages over the year (not at any one time). This assumes that the typical copier uses 874 kWh annually, and that a laser printer uses 283 kWh annual. These values are determined from a composite of different sizes of equipment.

6. Accounting for Losses Due to Auxiliary Equipment

As electricity passes through the UPSs and PDUs some is lost to the internal components in this equipment. With a full load, UPSs are approximately 95% efficient, and PDUs can be close to 98% efficient. As the load drops, however, these efficiencies decrease. Generally, 5% to 7% of the incoming power is lost as it passes through the UPSs and an additional 2% to 5% of the remaining power is lost to the PDUs.[51] Even under no load, however, there is a minimum amount of power needed.[52] Since these systems were generally under light loads, I assumed that the PDU and UPS efficiencies were on the lower end of this range and that the losses were approximately 5% and 7%, respectively. As a result, approximately 22 kW were needed to run the PDUs and 32 kW were needed for the UPSs, for a total of 54 kW.

Other auxiliary equipment such as building controls, fire alarms, security systems, telephone systems, and diesel generators also use small amounts of power. For the purpose of this calculation, I assume that these systems use roughly 2% of the total incoming power, or approximately 30 kW. Electrical line losses also require some additional power. Under a light load such as the load in this facility, line losses usually account for approximately 1% of incoming power. As a result, I assumed that line losses within the facility accounted for 15 kW.

Overall, I estimated that approximately 100 kW was used by auxiliary equipment and line losses. While these power draws occurred throughout the facility, I allocated this power to the active computer rooms since the majority of this auxiliary equipment was in the building for the sole purpose of supporting the computer room.

[50] Cited in Wilkins and Hosni, June 2000.
[51] Callsen, Thomas P, "The Art of Estimating Loads," Data Center Issue 2000.04, August 2000. Data also supported by discussion with MGE technical representative, 28 April 2001.
[52] The internal components include but are not be limited to inductors, capacitors, monitors, and filters.

7. Additional Power Needed For Cooling

The power consumed by computer equipment is totally converted to heat.[53] Mechanical equipment also consumes electricity and ultimately produces heat. In addition, the occupants of the building and the external environment produce some heat. Since the number of occupants in the building is small, and since the weather plays a small role in comparison to the computer equipment, I used the total number of watts determined in the steps above as an indicator of the heat load. Together, the heat load from the computer equipment and the other loads from the steps above equal approximately 911 kW. Because some of the HVAC equipment is rated in "power needed per ton of cooling," I converted the 911 kW heat load into "tons of cooling." One kilowatt is the equivalent of approximately 3415 British Thermal Units per hour (BTU/h), and 12,000 BTU/h is the equivalent of one ton of cooling. The 900 kW, therefore, is equal to approximately 260 tons of cooling. Some additional heat is released by the motors and fans associated with the HVAC system. The total heat load in this facility, as indicated by the monitor on the chiller, was approximately 320 tons.[54]

In order to estimate the amount of power required to operate the central plant, I estimated the power consumption of each of the components (chiller, cooling tower and pump). The active chiller in this facility was an 800 ton York chiller. (An additional 800 ton chiller was also on site as a backup.) The chiller required approximately 0.52 kW/ton, thus approximately 166 kW were needed to run the chiller.[55]

The active cooling tower had a 30 horsepower, approximately 22 kW, motor. However, since the cooling tower was running at only 40% of capacity, the motor was using the minimum amount of power: 2.2 kW or 10% of the design.[56]

[53] The Uptime Institute, 2000.
[54] Discussions with the facility manager, 02 May 2001. According to the facility manager, the heat load in this facility ranged from 300 tons in November to 350 tons in May 2000 when new clients were added. The 320 ton value is the best estimate for January 2001.
[55] The 0.52 kW/ton power requirement is from the manufacturer's information for this chiller.
[56] Data from manufacturer.

While the chiller and the cooling tower were operating at only 40% of capacity, the pump was used for a constant-flow water loop with a three-way valve that required a constant horsepower regardless of the load.[57] The pump, therefore, required full power or approximately 45 kW.[58] An alternative, more-efficient design would have allowed the pump to run at approximately 40% when the chiller was at 40% of capacity.

The central plant (including the chiller, cooling towers and pump), therefore, required approximately 213 kW. (See Table 7.)

Table 7. Central Plant Power Requirements

Central Plant	Power Required
Chiller	166 kW
Cooling Tower	2.2 kW
Pump	45 kW
Total	**213 kW**

In addition to the power to run the central plant, electricity is also needed to distribute the cool air throughout the building. Each of the computer rooms in this data center used six computer room air conditioning (CRAC) units to distribute cool air. Each month, five of the six units were operating while the sixth was down for routine maintenance. (This additional unit was for redundancy purposes.) Two of the computer rooms employed 50-ton CRAC units while one used 30-ton units. In the computer room with the smaller units, there were four additional air conditioning units located on the second floor that cooled the air remotely and then blew the cool air into the computer room. Overall, therefore, there were 22 units, 18 of which would usually run at one time.

Under the current light loads, these units were operating at approximately 30% of capacity. The fans within these units, however, ran constantly. The fans in a typical 50-

[57] From data center designer.
[58] From manufacturer's specification.

ton CRAC unit might require approximately 10 horsepower or 7.5 kW. The fans in the smaller 30-ton units would use slightly less power. Dehumidifiers and reheat systems as well as internal monitors and other components would add to the power requirements. A high-end CRAC unit with all of these features might require closer to 40 HP or approximately 30 kW. Assuming that 5 of the CRAC units were able to dehumidify and reheat and that the others were just cooling units, the 22 units would use a total of approximately 215 kW. In addition, the office space on the second floor has its own packaged air-handling unit. Additional fans were also necessary throughout the building. As a result, the total air distribution system for this facility could require close to 250 kW. (Note that this would mean an additional 70 tons of heat from the fans that would need to be cooled.)

The estimates above indicate that the total HVAC system in this facility, including the central plant and the fans for air distribution, used approximately 463 kW.

8. Calculating Total Power Needs

The power density assumptions for each part of the building are listed in Table 8. By multiplying the power density for each area by the appropriate floor area, I determined that this facility drew approximately 1.4 MW of power in January 2001. (See the shaded box in Table 9.)

The computer rooms in this facility were designed so that the computer equipment could draw an average of 60 watts of power per square foot (i.e., *design* computer power density = $60W/ft^2$). As shown in Table 8, however, the *actual* computer power density was less than 16 W/ft^2—just over one-fourth of the design value.

As mentioned earlier, however, this value is not representative of the total power needed to support the computer room. While the computers in this area drew approximately 16 W/ft^2, most of the additional systems in the building, such as the PDUs, the UPSs, and the back-up generators were in the building to support the computer room. In addition, a

Table 8. Breakdown of Power Density By End Use

Area Breakdown	Floor Area (ft^2)	Direct Use Power Densities (W/ft^2)			Supporting Equipment Power Densities (W/ft^2)			Power Density (W/ft^2)
		computers or prior use	lights	other	auxiliary equipment	central chiller plant	fans, CRAC units, AHUs	
Computer Rooms	27,500	15.7	1.1	0.0	3.6	4.4	7.5	32
Prior Use	12,600	20.0	N/A	N/A	N/A	2.4	N/A	22
Equipment Rooms	32,150	0.0	0.7	0.0	0.0	0.2	0.4	1
Office Space	14,300	0.1	1.8	0.3	0.0	0.3	0.5	3
Other Floor Area	38,775	0.0	1.0	0.1	0.0	0.4	0.7	2
Total Building	125,325	5.5	0.9	0.1	0.8	1.7	2.0	11

❖ Lights, other, auxiliary equipment and fans are for the "Prior Use" area are included in the 20 W/ft^2. Billing data for this area did not permit a more detailed breakdown.

Table 9. Total Power Demanded By End Use

Area Breakdown	Direct Use Power (kW)			Supporting Equipment Power (kW)			Total Power (kW)
	computer equipment or prior use	lights	other	auxiliary equipment/other	central chiller plant	fans, CRAC units, AHUs	
Computer Rooms	432	30	0	100	121	207	890
Prior Use	252	N/A	N/A	N/A	66	N/A	318
Equipment Rooms	0	23	0	0	6	10	38
Office Space	2	26	4	0	8	14	54
Other Floor Area	0	39	4	0	11	19	73
Total	686	117	8	100	213	250	1,374

large fraction of the energy used for the HVAC system was routed to cool the computer rooms and the auxiliary equipment. As a result, the actual power needed to support the computer rooms was much higher than 16 W/ft^2.

In order to determine the total computer room power density, which would include all of these support systems, I allocated the power used by the auxiliary equipment as well as the appropriate portion of HVAC power to this critical area. After including all of these support systems, I estimated that the total computer room power density for this building was closer to 32 W/ft^2. My key findings from this study are summarized in Table 10.

Table 10. Key Findings

Term	Definition	Results
Computer Power Density	Power drawn by the computer equipment (in watts) divided by the computer room floor area (in square feet)	16 W/ft^2
Total Computer Room Power Density	Power drawn by the computer equipment and all of the supporting equipment such as PDUs, UPSs, HVAC and lights (in watts) divided by the computer room floor area (in square feet)	32 W/ft^2
Building Power Density	Total power drawn by the building (in watts) divided by the total floor area of the building (in square feet)	11 W/ft^2

The majority of the power used to support the computer room went to the computer equipment. (See Figure 11.) The remaining power was used for the HVAC and auxiliary equipment as well as other end uses such as lighting. The HVAC system (including the central plant and the air distribution, or fans) accounted for approximately 37% percent of the power. Lighting represented only a small percentage—less than 3% of the power needs. These numbers indicate that targeting either the computers, or the HVAC system for energy efficient measures could potentially lead to large energy savings. (Further discussion of energy efficiency opportunities is included in Section VIII.)

Figure 11. Breakdown of Computer Room Power by End Use

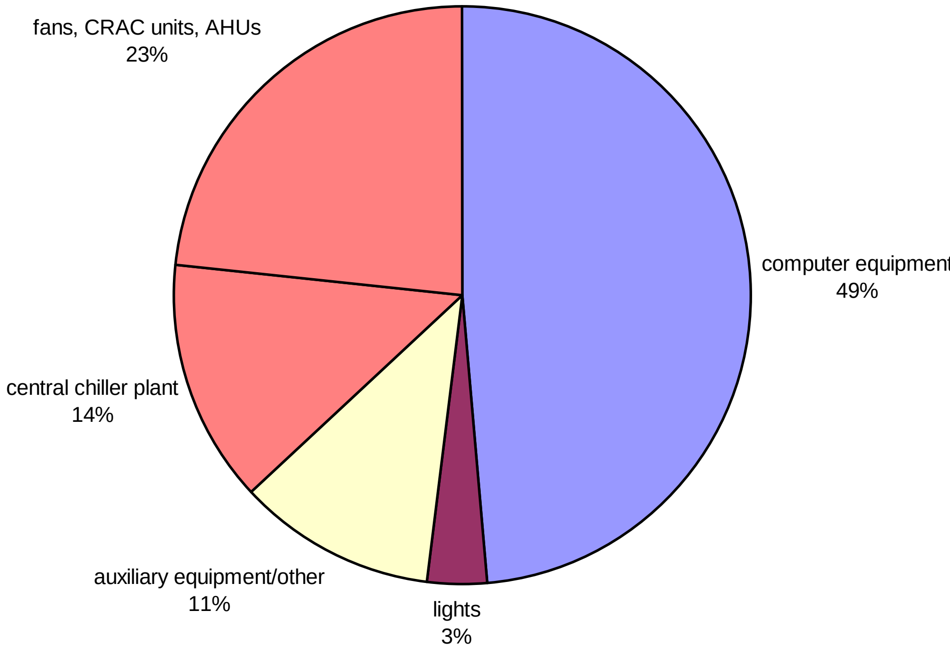

If additional computer equipment is added to the active computer rooms, the computer power density (32 W/ft^2) would increase, but not as much as one might guess. The support systems are already sized for full computer rooms, so additional diesel generators, PDUs or UPSs would not be needed to accommodate additional computer equipment. Furthermore, since the fans and the pumps do not have variable speed drive motors, they are already running at full power. Doubling the computer equipment, therefore, would not double the power requirements.

In the second stage of construction, an additional 19,000 ft^2 of computer rooms will be added on the second floor. While this will increase the building power density, it will not significantly alter the current total computer room power density. Additional computers, PDUs, UPSs, diesel generators, and other equipment *will* be added to support the new space. If the new computer rooms draw about the same amount of power as the current computer rooms (approximately 32 W/ft^2), then the new computer rooms will require an additional 600 kW, increasing the total facility demand to approximately 2 MW.[59]

VI. Using Measured Data to Confirm Power Needs

Billing data from this facility confirmed the estimates made in Section V. The Bay Area data center drew an average of 1.35 MW in January 2001. Electricity billing data also indicated that the average demand rose from approximately 500 kW to nearly 1400 kW between January 2000 and January 2001. See Figure 12. Billing data for additional months in 2001 were not available, but a final visit to the facility in April of 2001 revealed that two of the three computer rooms were still drawing about the same amount of power as in January, while the third had dropped by approximately 40 kW because of the loss of a major customer.[60]

Figure 12. Average Power Demanded by a 125,000 ft^2 Bay Area Data Center

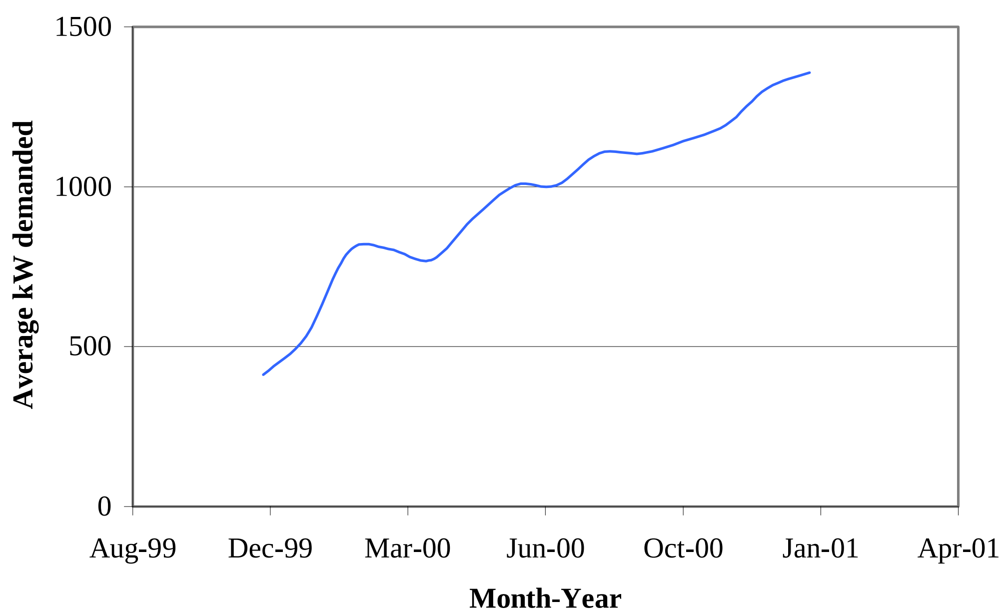

Source: Graphed from billing data, December 1999 to January 2001.

In order to confirm my findings, I reviewed billing data for four other data centers across the country. From these billing data and from information about the computer room area, I was able to determine an upper limit of the total computer room power density for these four facilities. (See Table 11.) I used the billing data to find average demand in the month with the highest consumption. (This was usually the most recent month.) I then divided the highest average power demand for the facility by the computer room floor area. This estimate of total computer room power density is an overestimate because it assumes that all of the power for the entire facility is used for the computer room. Even these overestimates, however, indicate that the total computer room power density is always less than 40 W/ft^2.

Table 11. Comparison of Five Data Centers

Location	Units	Bay Area Data Center	Data Center A	Data Center B	Data Center C	Data Center D
Building Area	ft^2	125,325	115,000	154,158	NA	358,362
Computer Room Area	ft^2	27,500	40,000	45,000	48,186	38,500
Building Power Density	W/ft^2	11	3	10	NA	4
Upper Limit for Total Computer Room Power Density	W/ft^2	32	8	34	38	35

- For the later four data centers listed, the Upper Limit for Total Computer Room Power Density was calculated by dividing the average power demand for the entire facility (from billing data) by the computer room area. This number includes all of the power used by the entire building and is therefore an overestimate. The Bay Area number is based on measurements in Section V.

According to the *Los Angeles Times*, a recent study by PG&E also found that the computer rooms in several server farms in PG&E's territory consume about 40 W/ft^2.[61] Unfortunately, however, this study is unavailable to the public and the article does not indicate how these measurements were taken, or whether this 40 W/ft^2 includes cooling or not. Additionally, an article in the January 2001 edition of *Network Magazine* indicates

[60] The average demand in one computer room went up by 2 kW and the second computer room went up by 9 kW.

[61] Reiterman, Tim, "San Franciscans Protest as 'Server Farms' Sprout," *Los Angeles Times*, 26 March 2001.

that EPRI members also have private measurements of power demands.[62] The article quotes Steve Rosenstock of EPRI as saying, "…Our members have been measuring what they're actually using after installation. It's closer to 25 to 40 watts per square foot.[63]" Again, however, the measurements are not publicly available, nor is it clearly stated how the measurements were taken or over what area the stated power density applies or whether these estimates include cooling.

Based on my findings in Section V, at low loads, the computer power density is roughly half of the total computer room power density. This would imply that the computer power densities in these facilities are all less than 20 W/ft^2. All of these facilities, however, have been designed to accommodate computer power densities that are between 60 and 90 W/ft^2. The actual power drawn by the computers, therefore, is less than a third of designed computer power density.

In addition to the billing data mentioned above, the facility manager at a New York data center had several measurements from January 2001 to the present (April 1, 2001). The available data included weekly readings for all PDUs, UPSs, and automatic transfer switches.[64] Demand at this data center was approximately 1.7 MW. Less than a quarter of this—less than 400 kW—was used to directly power the computer equipment.

I determined the computer power density in this facility by dividing power exiting the PDU by the computer room area—similar to the method performed in step 1 of Section V. While this data center was designed for computer equipment that draws 90W/ft^2, the actual computer power density was approximately 7 or 8 W/ft^2.

Based on the measurements taken at the New York facility, an additional 50 to 75 kW was used to run the UPSs and PDUs in this facility. Demand from the chiller plant at this data center was over 500 kW. The remaining 700 kW went to air distribution, lighting,

[62] Angel, Jonathan, "Energy Consumption and the New Economy," *Network Magazine*, 1 January 2001.
[63] Angel, 1 January 2001.

diesel generators, fire suppression and security systems, line losses and other support systems. Again, for this data center, even an overestimate of the total computer room power density indicates that the computers and all of the support systems drew less than 40 W/ft^2.

Based on the data from the five data centers that I reviewed, my best estimate of average total computer room power density is approximately 40 W/ft^2. In order to take into account the fact that there may be data centers that contain more computer equipment than the data centers that I studied, in Section VIII, I take a conservative estimate and assume that the average data center has a total computer room power density of 50 W/ft^2. This estimate includes all power consumed by the computer equipment as well as the support systems. This estimate is much higher than the average power density in an office building, approximately 5 to 10 W/ft^2. It is, however, much lower than assumed by many of the current estimates available publicly. While this estimate is still rough and may be skewed by the data available, it gives a ballpark estimate of the true power consumed by data centers.

[64] The automatic transfer switch is used to transfer the load from the AC grid to the diesel generator when a power outage occurs. The monitor on the automatic transfer switch, therefore, gives a reading of the total facility power load.

VII. Reasons for Exaggerated Forecasts

As noted earlier, data center power needs are sometimes overestimated because of inaccurate calculations. Often, the total building power load is estimated by multiplying the computer power density by the area of the entire building. For the data center described in Section V, the simple (but inaccurate calculation) of multiplying the design computer power density by the entire footprint of the facility (i.e., 60 W/ft^2 by 125,000 ft^2) would give an incorrect estimate of 7.5 MW—more than five times what this facility currently draws, and more than three times what this facility would draw even if the second floor were completed and occupied.

Aside from these reasons why the media or the general public may be overestimating the power needs of data centers, there are several more technical reasons why data center loads are often overstated even by engineers. Below I outline nine additional areas where data center design and the accompanying assumptions can lead to overestimates of power needs. While not all of these are a concern at all data centers, most data center estimates of power loads—and thus the requests they make to utilities—include some combination of the assumptions listed below.

1. The use of nameplate power consumption

All computer equipment is given a value for the theoretical maximum amount of power that the equipment can draw. This is also referred to as the nameplate value since it can usually be found engraved on the back or the side of the equipment. This is often the only estimate of power draw available for a piece of equipment. This value, however, is always overstated for the following reasons:

- Most devices use less electricity while running than they do at their peak, which is often during start-up. In fact, for safety reasons, most computer equipment never draws more than 80% of the rated power even at its peak.[65]

[65] Nordman, Bruce, LBNL memo, 5 December 2000.

- Devices like computers, routers and switches can have slots for add-in cards. The power rating must be sized as if all slots were full with cards that draw the maximum amount of power. If all of the slots are not full, or if the slots are designed for more than the cards that are in use demand, the power requirements of the equipment will be overstated.
- For convenience, manufacturers standardize power supplies across multiple product lines in order to minimize the number of different power supplies that they have to produce. For some equipment, therefore, this would lead to larger power supplies than required even given the information above.
- Finally, power supplies are often oversized in anticipation of future upgrades.

A Swiss report by Basler and Hofman recorded measurements of the electricity consumption of network components (routers, switches, multiplexers, micro repeaters, media converters) in two modern networks with 82 and 1200 users respectively.[66] Basler and Hofman found that the measured power was approximately 30% of the nameplate specifications.[67]

This would mean, for example, that a switch in the U.S. might have a nameplate rating of 16A or 1.9 kW but would only draw 4.8A. In this example, the supply, and thus the amount of power requested, is overestimated by 70%. LBNL has measured several pieces of computer equipment similar to the type found in data centers and has confirmed that nameplate does overstate actual power demands. A Sun Ultra server, for example, that was rated to use 4A, or approximately 475 W, was measured to use approximately 113 W when running normally, and a maximum of 142 W at startup.[68]

[66] Basler and Hofman, 26 November 1997.
[67] Basler and Hofman, 26 November 1997.
[68] Nordman, Bruce, LBNL, data on server measurements taken at Soda Hall, UC Berkeley in January 2000, email, received 17 January 2001.

2. Oversized circuits

Given that the nameplate rating usually overstates actual power demands, a customer may decide to plug four 6A servers into a 20A circuit. Most likely, however, due to the desire for redundancy and secure power, a customer would not risk the possibility of a power loss due to an overloaded circuit. Thus, if the three servers were rated at 6A each, and the circuit was a 20A circuit, the customer would most likely plug a maximum of 3 servers, or 18A, into the circuit and leaving an additional 2A of infrastructure free. In some data centers, there may be even more un-utilitized capacity. At HostPro, a fact sheet explains that, "To conform to electrical code for peak power use, maximum power usage is limited to 75% of circuit values (e.g. 15 amps for a 20 amp circuit). HostPro reserves the right to audit customer circuits at random to verify power usage.[69]" Trying to size more accurately would require closer monitoring, but at large data centers, "it is extremely difficult to monitor every circuit for individual usage because there are many many thousands of circuits.[70]" This unused infrastructure can add up. For a rack with nine 6A servers, a customer would request three 20A circuits. This would mean that the infrastructure would be oversized by more than 10% (6A out of 60A), or in the case of HostPro, it would always be oversized by at least 33%. Since the co-location facilities are built long before the mix of internal equipment is determined, it is difficult to minimize the oversizing of circuits and other infrastructure. A facility is not likely to rewire circuits for incoming customers.

3. Dual power feeds

Some computer equipment employs dual power supplies in order to ensure that the computers do not lose power in case of a power supply failure. In this case, even though the equipment might draw a maximum of 6A, it would have two 6A power supplies. Most equipment tends to draw only slightly more power with two power supplies than with one, but all dual power supply equipment is designed so that it could run entirely off

[69] HostPro, data center fact sheet, www.hostpro.com, viewed 4/11/01
[70] Anonymous, communication with Exodus employee, 19 March 2001. Points 1-3 in this section build on Bruce Nordman's earlier work.

of one power supply if the second fails.[71] Each power supply would be plugged into a separate circuit. In this case, to run the same three 6A servers mentioned above, you would need two 20A circuits—approximately 100% more than the equipment would draw even if it required the nameplate power.

4. Overestimates of equipment in the rack

Assuming that the majority of racks available are approximately 6 feet tall, it is possible to estimate the number of computers that can fit into these racks. This is not, however, as easy as it may seem. A couple of years ago when the newest data centers were being designed, the majority of severs were approximately 7 inches tall, or 4U. A 6-foot rack, therefore, could hold approximately 10 servers. Today, many servers are only 1.75 inches tall, or 1U, so a single rack can hold approximately 40 servers. The electricity use, however, has not declined at the same rate as technological compaction. The energy use of a server is based on how many processors and drives it has. Since today's 1U servers can have as many processors as a 4U server, the 1U server might consume about the same amount of electricity. In order to provide customers with the desired power, estimates are often determined based on the assumption that these racks could each hold 40 servers. Most data centers, however, still use at least some larger pieces of equipment.

5. Facilities that are not full

Regardless of how many pieces of equipment could fit in a rack, many racks are not fully utilized. In the Bay Area data center described in Section V, 47% of the audited racks had no electrical equipment, and many others were not filled to capacity. The average rack was only one-third filled. As a result, the infrastructure, and the estimates, assume more racks than are used, again leading to overestimated loads.

[71] Bruce Nordman, "Electricity requirements for LBNL's Networking Hardware," memo to Jon Koomey at LBNL, LBNL Publication 835, 9 December 1999.

Furthermore, while revenues or payback periods are usually calculated based on the data center being filled to 30-40% of capacity, power requirements assume that the facility will be fully utilized.[72] Data centers may give the utility a build-out schedule, but usually the initial requests for power err on the high side in order to ensure that they have the power needed.

It is unlikely, however, that these data centers will ever be filled to capacity. The racks in the three data centers mentioned in Section V above were all approximately one-third full—with the managed hosting areas slightly more full than the co-location areas. It is difficult to predict use rates in co-location facilities since the customers are not known ahead of time. In addition, outside companies may also build in redundancy by renting additional unused space in case they need it. In some data centers, companies pay a "reservation fee" to reserve racks or cages that they may never use. This is usually because the companies are anticipating growth, but it is too early to know how many of these racks will ever be filled.

6. Estimates based on anticipated loads

As mentioned above, servers have become much more compact over the past couple of years. A recent paper by the Uptime Institute, using information from 15 computer manufacturers, shows the historical trend (from 1992 to present) of power used by a rack of servers. The graph from the Uptime Institute indicates that a full rack of servers today use about 600 to 900 W/ft^2 (where the footprint is the footprint of the rack, or approximately 6 ft^2.) This paper also indicates that the same rack of servers could require 1200-1700 W/ft^2 by 2005.[73] Given the rapid introduction of 1U servers, and the rapid turnover of computer equipment, data centers have started designing for the future. What the future holds, however, is unknown. Computer electrical loads are likely to increase at a much slower pace in the years to come for several reasons. Computer chips, for example, will most likely be much more efficient. In fact, "Four companies recently

[72] Mahedy et al., 3 August 2000.
[73] The Uptime Institute, 2000.

announced plans to release Internet servers with low-power chips.[74]" Although hard to quantify, building for the future also leads to overestimated requests for power.

7. Overestimated HVAC

Overestimating the power needs of the computer equipment means overestimating the heat load that will need to be cooled. A recent paper in the *ASHRAE Journal* is indicative of the fact that HVAC systems are often oversized. The study measured the actual cooling needs of office equipment and clearly determined that the nameplate rating overstates the actual cooling needs and "should be ignored when performing cooling load calculations.[75]" This problem is even more prevalent in a data center where computer equipment density is much higher than in office buildings. Sizing the HVAC system to an overestimated load will require larger chillers and fans, and more computer room air conditioning (CRAC) units than needed. All of these things require power. The electrical system will have to be sized to accommodate a fully running HVAC system despite the fact that some of this mechanical equipment may not run or will be used only at partial capacity if the heat load is not as high as expected. Thus, under the assumption that all of the mechanical equipment is operating, the estimate of the facility's power requirements becomes even larger.

8. General overdesign and safety factors

Engineers also typically build in safety factors. In an industry where reliability is highly valued, and the engineers know that they will be blamed if the system crashes, it is likely that several systems will be overdesigned. For example, the mechanical system may be oversized by 20%. It is also possible that additional safety factors will have to be incorporated to account for difficulties in balancing loads. In data centers, it is common for the equipment on the three phases of the incoming power supply to be unbalanced. As a result, electrical engineers usually oversize the neutral bus to account for mismatched

[74] Stein, Jay, "More Computing Power, Less Electrical Power," in *ET Currents* Number 7, E Source, March 2001.
[75] Wilkins and Hosni, June 2000.

loads. While this is common, mismatched loads can also mean that a PDU or UPS would only be able to support 70 to 75% of its capacity; thus a data center would need to add additional auxiliary equipment. Additional equipment means even larger power demands.

The oversizing of each system is further compounded by the fact that the engineers that design the mechanical systems are not the same engineers that design the electrical systems or the IT equipment. Each discipline adds its own safety factors. The electrical system, therefore, will be oversized for an already oversized IT and mechanical load.

9. Overestimates of the number of data centers

Finally, overall estimates of the number of data centers that will be built may also be overstated. Several companies may not end up building the data centers that they originally planned. Companies that have put in requests for power for 2003 may be out of business long before then. It is also possible that several of the speculative data centers are being double counted. Data center owners that are planning to build a new facility may go to more than one utility with requests for power because they are "shopping." Data centers understand that utilities operate on a longer time schedule. It often takes the utility longer to respond to requests for new substations or transformers than it takes for the data center to be built; so companies enter requests in several areas even though they have not selected a final site. The speculative power requests can also lead to overestimates.

In sum, assumptions about the power needs of data centers are based on design criteria that incorporate oversized, redundant systems, with several built-in safety factors. In addition, the values that are commonly cited in the media as well as in discussions with utilities assume that the data centers will be filled to capacity. The estimates of power used by these facilities, therefore, are likely to be greatly overstated.

VII. Implications of Findings

Nationwide, approximately 9.5 million square feet of data center space was devoted to computer equipment in 2000. If the average data center has a total computer room power density of 50 W/ft^2, then across the country these facilities would require less than 500 MW of power, and would use a fraction of a percent of all electricity used nationwide. (See Table 12.)

Table 12. Nationwide Electricity Demands from Data Center Hosting Facilities

	Units	2000	2003 low	2003 mid	2003 high
Computer room floor area	Million ft^2	9.5	20	25	30
Total computer room power density	W/ft^2	50	35	60	85
Data center total power	MW	475	700	1500	2550
US electricity use	TWh	3364	3608	3608	3608
Data center electricity use	TWh	4	6	13	22
Data centers as % total electricity use	%	0.12%	0.17%	0.36%	0.62%

- Total U.S. electricity use from Energy Information Administration's *Annual Energy Outlook 2001*.

Inevitably, the number of data centers—as well as the net area of these facilities—will rise. In addition, it is possible that the total computer room power density will increase. As the table above shows, however, even with high (yet still realistic) estimates, by 2003, data center power requests will add up to less than one percent of all electricity consumption nationwide. Moreover, it is important to note that some unknown portion of this demand is not actually new electricity demand. Some of the computers in these data centers are just relocated from corporate office buildings to data center hosting facilities.

While they represent only a fraction of the total electricity consumed in the United States, the electricity demands from data centers are significant in certain locations. As a result, it is possible that local distribution problems will occur in regions where data centers are concentrated. In the Bay Area, for instance, the Salomon Smith Barney report estimated that there were 1 million net square feet of data center space at the end of 2000.[76] Using the assumption of $50W/ft^2$, data centers could require 50 MW of power in the Bay Area alone—approximately 10% of the total demanded by data centers nationwide. In the Bay Area, this could mean approximately 438 GWh of electricity a year, or approximately 1.2% of electricity consumption in this area.[77] By 2003, power demands from Bay Area data centers will probably increase.

Figure 13. Regional Data Centers Power Needs

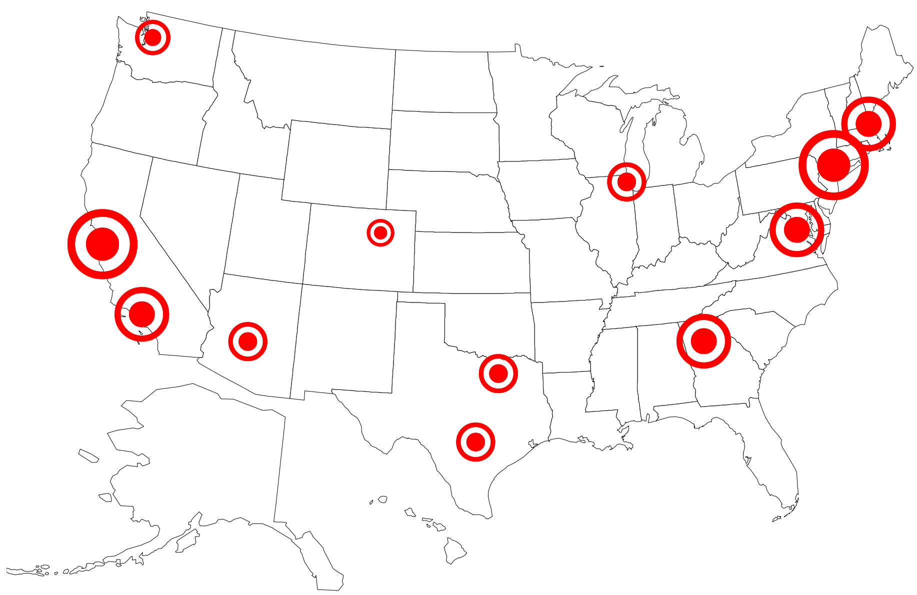

❖ The markings are indicative of the relative power needs of data centers in these regions. Based on Salomon Smith Barney data (from Mahedy et al., 3 August 2000) and using a conservative assumption that the average total computer room power density is $50\ W/ft^2$.

Data centers appear to be ideal electricity customers: they demand a relatively steady amount of power 24 hours a day. In reality, however, even after these facilities are built, utilities do not always know the true demand since the industry is just starting to grow and changes are occurring rapidly. For a utility, not having a good sense of the data

[76] This estimate does not include the 2.2 million (gross) square foot U.S. Dataport facility.
[77] According to the CEC data cited in Table 2, total Bay Area electricity use is approximately 35,400 GWh per year.

center's electricity demand leads to difficulties in providing supply and hedging against risks.

Utilities also face the challenge of meeting the customer's demand for infrastructure. This is especially difficult given the distinct differences in timing and planning cycles between utilities and the Internet industry. Utilities are accustomed to getting two or three years notice for new large office buildings and industrial centers. Now, they face the challenges of putting in power lines, transformers, and substations within a few months.

Currently, utilities face a "lose, lose, lose" situation. If the requests they face are accurate, utilities will be unable to meet the need in the short time frame. If the requests are indicative of future demand and utilities decide to upgrade incrementally, constantly resizing the infrastructure to meet the data center's needs will be time consuming and costly. And if the requests are overstated and never materialize, utilities will overinvest in infrastructure and will not be able to recover their costs. To utilities, accurate estimates of power needs are extremely valuable.

For most data centers, however, energy costs are not high on the list of priorities. Data centers rent space for around \$200-\$500/ft^2 per month.[78] For a 40,000 ft^2 facility with 20,000 ft^2 of core data center space, if half of this space is rented out, this could mean revenues of approximately \$3.5 million a month. Electricity bills for this same facility, however, might be on the order of \$72,000 per month, or just 2% of income.[79] The small outlay for energy, as well as the fact that the bills usually go to the corporate office rather than the data center facility manager, mean that good estimates are not usually available. Furthermore, until recently, most data centers paid only for the energy that they use—not for what they estimate in their initial proposals.

[78] Mahedy et al., 3 August 2000. Supported by anonymous source with Bay Area Internet company that rents data center space, 12 March 2001. This may be changing with the recent economic downturn.
[79] Mahedy et al., 3 August 2000. Estimate based on \$.0.10 per kWh (which would include an average demand charge.)

In PG&E's territory, for example, data centers that fall under schedule E-19—commercial and industrial users that consume more than 499 kW—pay a customer charge ($175.00/meter), a demand charge (per max kW in a period) and an energy charge (per kWh). There is no charge for the cost of installation or for the ratio of a load's actual energy consumption over a period of time to the maximum amount requested. Since the cost of building the supply infrastructure is related to the maximum amount of power requested (i.e., the capacity of generators and transmission lines) whereas the revenues from electricity sales are related to the amount of energy (kilowatt-hours) consumed, the result could be large, uncovered (or stranded) costs. For example, although data centers requested 341 MW last year in PG&E's territory, based on the estimates above, it is likely that less than 50 MW was needed to serve these customers.[80]

In order to avoid excessive risk and act in the best interest of shareholders, several utilities have started to charge data centers based on their initial requests for power. Last year, ComEd started charging "server farm deposits range from $500,000 to more than $10 million per project, depending on the amount of engineering work, equipment and installation needed.[81]" "The utility's contract pays developers' deposits back in portions over a five-year period depending on how a site's electricity usage progresses towards the original load estimate.[82]"

Seattle has also started to implement a similar rate tariff: "Puget Sound Energy…asked the Washington Utilities and Transportation Commission to accept a tariff on new data centers. The tariff is designed to protect the company's existing customers from footing the bill for new base stations necessary to support the projects. Those stations could cost as much as $20 million each.[83]" Four other utilities: Nstar in Boston, Consolidated Edison in New York, Commonwealth Edison in Chicago, and Southern California Edison are also considering a similar charge. According to Source One, a utility aggregator, **"**The

[80] The 341 MW value is from Energy Solutions and Supersymmetry, 19 October 2000. The 50 MW is based on my findings.
[81] Ahlberg, Erik, "Electricity Utilities Fear Drought at Server Farms," *Dow Jones*, 30 January 2001.
[82] Ahlberg, 30 January 2001.
[83] Cook, 5 September 2000.

utilities are taking the position that if it's not real, then at least [the hosting companies should pay] for the capital improvements.[84]"

Utilities are also looking into other alternatives such as assistance with data center design. California's three largest utilities run a program called Savings By Design which provides the owner of the facility and the design team with energy design tools and information to improve building performance.[85] The program also provides financial incentives to design teams that are able to improve the energy efficiency of buildings. While this program is not specifically set up to assist data centers, a similar program could encourage energy savings in these facilities. Working with data center designers is a creative and effective approach because the same design firms are contracted to work on several data center facilities. While each building is unique, a significant amount of knowledge about the design of these facilities is transferable between projects. Although technological advances will continue to require adaptability, the ability to transfer knowledge and learn between projects will improve data center designs. Given the right tools and incentives, designers can think through current and future phases of the project and design flexible modular systems that will lead to the most efficient buildings.

Current energy efficiency options include raising temperature set points and switching to more efficient HVAC systems. Several HVAC design engineers that I spoke with indicated that the power requirements of air conditioning equipment in data centers are much greater than anticipated based on the computer load due to built-in redundancy and inefficiencies of the cooling system. The redundancy in the HVAC system depends on the needed reliability of the data center. Typically, large data centers require 25% redundancy. There are, however, HVAC systems in data centers with as much as 400% redundancy installed.[86] Poorly designed HVAC systems in data centers use at least twice the electricity as a more efficient system.[87]

---

[84] Qualters, Sheri, "Energy costs surge for new projects," *Boston Business Journal*, January 29, 2001.
[85] PG&E website, www.pge.com/003_save_energy/003b_bus/003blcl_program_info.shtml, viewed 12 April 2001.
[86] ASHRAE/ATCE, information from an anonymous ASHRAE HVAC engineer, 13 March 2001.
[87] ASHRAE/ATCE, 13 March 2001.

The power requirement of the fans needed for cooling a building can often be estimated by using a standard equation. Using the standard equation, (which assumes a central air handling unit rather than the distributed CRAC units), the additional energy required to distribute cool air to this facility should be roughly one-third of the energy needed by the CRAC units.[88] Thus, in the facility that I studied, the electricity demands may be able to be reduced by 150 MW by using an alternative air distribution system. For most data centers, the standard calculation significantly underestimates the true power used to distribute air because CRAC units are inefficient. Thus, there is room for energy efficient improvements in the air distribution system. Furthermore, while this data center employed a chilled water central plant (usually 0.5 to 0.7 kW/ton), many data centers still use less efficient direct expansion air cooled systems (1.2 kW/ton) because they are worried about the proximity of water to computer equipment. In data centers with direct expansion air-cooled systems, therefore, there is room for additional energy efficiency improvements.

These facilities also have several other design inefficiencies. The current electrical design converts incoming electricity from AC to DC to AC to DC before being used by the computers. These conversions lead to excessive power losses that could be eliminated. Doing so, however, would require redesigning all Internet computer equipment so that DC powered equipment was standard. While not impossible, a move toward redesigning computers would require a large scale cooperative effort by several parties.

Since the computer equipment accounts for the majority of the power requirements, it is also important to focus efforts on increasing the energy efficiency of this equipment.

[88] Assuming that approximately 500 cubic feet of air per minute would be circulated for every ton of heat, if the static pressure (SP), and the efficiency of the fans, motors and drives, are known, it is possible to estimate the power needed to distribute air by plugging these numbers into a standard equation as follows:

$$\text{FAN POWER (kW)} = \left[\left(\frac{\text{500 cfm air}}{\text{ton heat}} * \text{320 tons heat} * \text{3 inches SP}\right) \div \left(\frac{\text{6,300 inches SP - cfm air}}{\text{HP}} * \text{0.7 fan - motor efficiency}\right)\right] * \frac{\text{0.746 kW}}{\text{HP}}$$

Computer manufacturers are already developing energy efficient computer equipment. Chip manufacturers are designing new chips that require as little as one-tenth of the power of their predecessors.[89] Whether this computer equipment will be adopted remains to be seen. Companies in this industry are hesitant to try new approaches that may threaten reliability. Utilities and industry groups such as the Silicon Valley Manufacturing Group can, however, take steps to encourage and promote the use of these new technologies.

Alternatively, discussions with several utility representatives, designers, and industry professionals have led to the suggestion of a data center "smart park." Data centers could be located in an industrial park that would provide reliable and redundant power from a combined-cycle natural gas facility located onsite, cooling from an energy efficient chilled water system, and other built-in energy efficient systems.[90] U.S. Dataport has proposed three large (2.2 million to 3.7 million gross square feet) data center facilities in San Jose, New York, and Northern Virginia that, while not as efficient as the proposed smart park, could be guided toward an energy efficient design. The San Jose complex, however, has been met with resistance. Most likely, if this complex is built, it will have to provide its own power. If done correctly, this type of facility could encourage clean sources of power and energy efficient designs. An energy efficient data center complex might be an appealing solution for several areas, such as Sacramento and San Francisco, that have begun passing local zoning ordinances to limit data centers.[91] According to Peter Fortenbaugh, Senior Vice President of strategic planning for Exodus Communications, data centers may be evolving toward this sort of utility model.[92] However, without strict zoning regulations and a cooperative effort to make sure that these types of facilities are built to encourage efficiency, the smart park idea will not be successful.

---

[89] Stein, March 2001.
[90] Austin Energy Conference, 13 February 2001.
[91] McCarthy, Mike, "Ordinance would limit downtown 'telecom hotels,'" *Sacramento Business Journal*, 21 July 2000; and "Council Oks restrictions on downtown 'telcom hotels,'" *Sacramento Business Journal*, 31 August 2000.

Whether the future brings computer equipment with even larger power requirements is unknown. It is certain, however, that this industry is young, and that there are several opportunities to encourage more energy efficient measures in order to shape future data centers. There are energy efficient solutions that will enable data centers to be built and to support local economies without draining local power supplies or damaging local environments. It will, however, take the cooperative efforts of several stakeholders such as utilities, industry groups and local governments. In order to take these steps, it is important for these groups to understand the true power demands. Thus, utilities and data center companies need to monitor the electricity consumption of these types of facilities and make this aggregate information public.

[92] Schmelling, Sarah, "The Frugal Data Centers," *the Net Economy*, 2 April 2001.

IX. Conclusions

While some of the discussion in this paper may seem simplistic, not understanding the metrics or not clearly stating what is being discussed can lead to enormous amounts of confusion. In order to arrive at more accurate assumptions, it is critical that people explicitly state whether the power density that they are discussing is a design criteria or an actual measured value, whether it includes all power uses or just the computers, and whether it is for the entire building floor area or just the computer room. Misinformed forecasts can lead to inaccurate calculations. Moreover, even the more accurate

Figure 14. Which Numbers Should You Use To Determine New Growth Due to Data Centers?

An Example From This Study

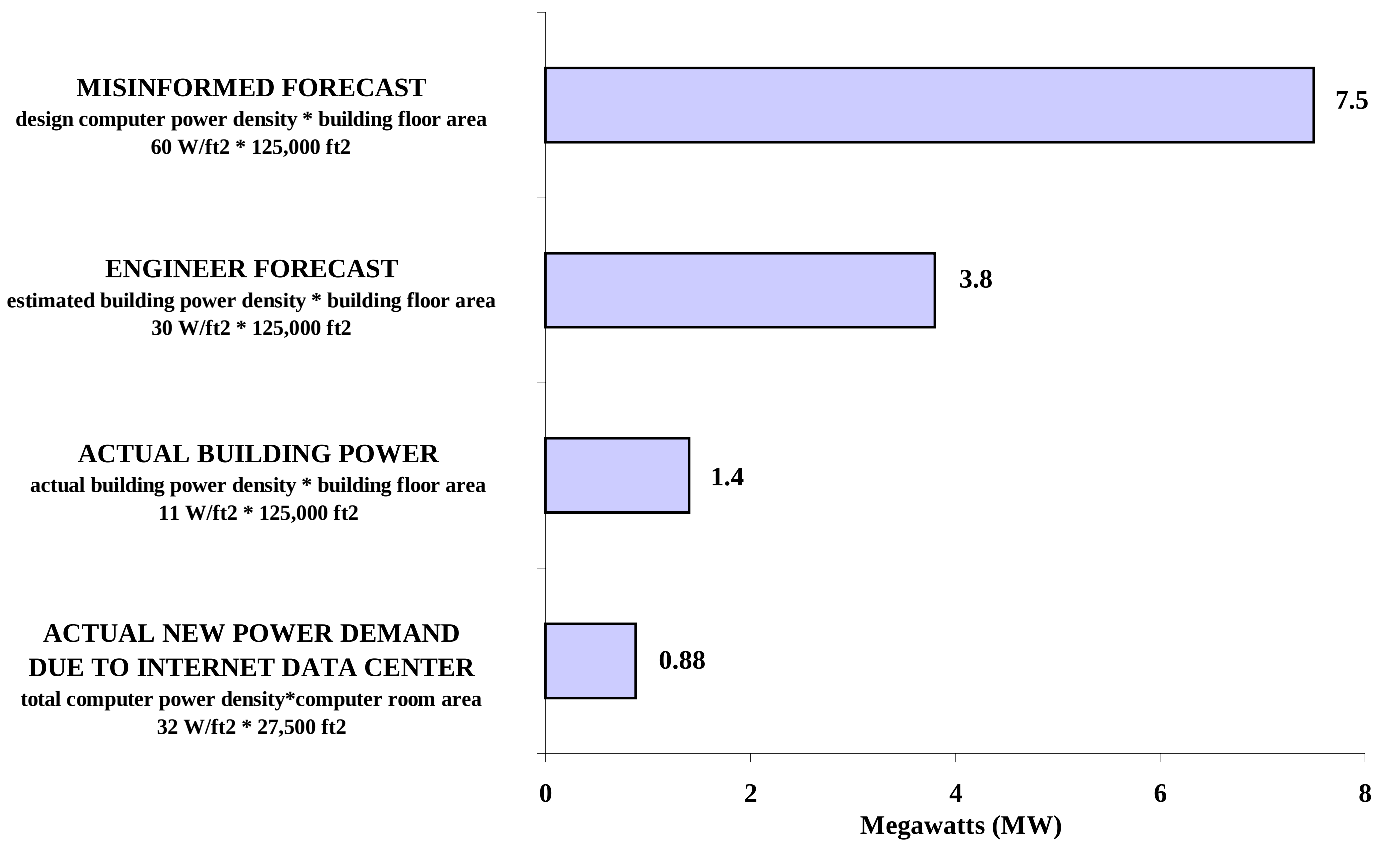

- These numbers are based on the Bay Area data center that I studied.
- Since this is a renovated multipurpose building (with some office space and some space remaining in its prior use) a portion of the building power is not new demand due to the data center. Only 880 kW is representative of new power demands by the data center.

"engineering" assumptions may still significantly overstate new power requirements from data centers. Figure 14, above, uses information from one specific data center to give a schematic representation of the overestimates that can occur. These numbers are for the Bay Area data center discussed in Section V. The first step toward developing more accurate estimates and projections is to use clear terminology and common metrics in order to eliminate confusion.

Based on my findings, the electricity requirements of this industry do not translate into a national crisis. Even high estimates of power densities indicate that demands from these facilities will require less than one percent of U.S. electricity consumption or only 22 TWh per year by 2003. However, power requirements in data centers are much larger per unit of floor area than the requirements of a commercial office building. Therefore, there is room for energy efficiency gains in current facilities as well as in data centers that will be built in the future. Lower power servers and better-designed HVAC systems, in particular, offer options for significant energy savings. Energy efficiency improvements will help to reduce local impacts that may occur in data center hubs.

One of the largest difficulties will be trying to bridge the gap between energy efficiency and the reliability/redundancy requirements of this industry. For most data centers, reliability is so important that data center owners and designers rely on proven methods and do not want to test new energy efficient options. Furthermore, most data centers are seeking to minimize the time it takes to enter the market.

It is important that utilities understand the factors that may lead to overestimates and that they provide the right incentives to encourage more accurate estimates. In addition, utilities can play a key role in encouraging energy efficient solutions since they already have established relationships with data centers.

Computer manufacturers and local governments can also play an important role in finding ways to reduce energy consumption in data centers. New low-power servers and zoning regulations are helping to push this industry in the right direction.

Most importantly, a collaborative effort between utilities, data centers, and local governments is needed to better understand the real power needs of data centers. Detailed studies of current energy requirements will help all parties to understand the needs of this industry and will help to provide insights into where energy efficiency measures can be the most effective.

Acknowledgements: Special thanks to Jon Koomey, whose guidance (and recently released book) helped me to "Turn…[my] Numbers into Knowledge." Thanks also to Michele Blazek, Bruce Nordman for their help and support of my project. I couldn't have done this without you. I also want to thank my other reviewers including Dan Kammen (a jack of all trades) and the LBNL researchers who gave this report a final read through, as well as the designers and engineers (especially Jim Morris and Steve Greenberg) who assisted me and answered numerous phone calls and questions. Of course, this project would not be possible without the help of the companies that opened their doors to my research, and the employees at the data center who led me through the facility and put up with my repeated requests for information. Unfortunately, for confidentiality reasons, I can't mention them by name, but a sincere thanks to all of you. I really appreciate your help. And to the facility manager from NY who had been recording data at his facility, you've done an incredible job. Hopefully others will follow in your footsteps. Lastly, but most importantly, thank you Phelps. In a nutshell, you're the apple of my eye.

Appendix A: Detailed PDU Data

PDU Data Collected By Jennifer Mitchell-Jackson, January 2001

Data Center Area	PDU #	Average Volts	Phase A	Phase B	Phase C	Neutral	
One	1a	0	0	0	0	0	
One	1b	276	16	18	18	0	
One	2a	276	23	26	25	0	
One	2b	0	0	0	0	0	
One	3a	277	40	41	46	0	
One	3b	0	0	0	0	0	
One	4a	0	0	0	0	0	
One	4b	276	49	43	52	0	
					Total Amps	397	Amps
					Power (in*)	110	kVA
					Power (out)	101	kVA
Two	1	120	70	53	69	0	
Two	2	120	56	54	55	0	
Two	3	119	66	60	49	46	
Two	4	119	74	93	72	63	
Two	5	119	73	60	95	50	
Two	6	120	128	100	64	99	
Two	7	120	80	61	29	58	
Two	8	120	60	32	37	44	
					Total Amps	1590	Amps
					Power (in)	205	kVA
					Power (out)	190	kVA
Three	1a	0	0	0	0	0	
Three	1b	276	66	0	65	60	
Three	2a	0	0	0	0	0	
Three	2b	272	66	55	69	0	
Three	3a	0	0	0	0	0	
Three	3b	275	42	45	47	0	
Three	4a	0	0	0	0	0	
Three	4b	276	41	40	42	0	
					Total Amps	578	Amps
					Power (in)	159	kVA
					Power (out)	146	kVA
				All 3 Data Centers (in)		474	kVA
				All 3 Data Centers (out)		437	kVA

*N.B. “in” means in to the PDU and “out” means out of the PDU.

Appendix A continued.

The PDUs in the newer computer rooms were designed to display ***input*** amps and volts at 480/277V, while the older PDUs in the middle data centers displayed ***output*** amps and volts at 208/120V. To convert input to output (and vice versa) I assumed approximately 5% losses due to the transformer and other internal components in the PDU. Apparent power was converted to real power using a power factor of 0.97 as described in the text.

Summary Table of Measured Input and Output Data

Computer Room	Average Volts	Total Amps	Apparent Power Consumed By Computers	Real Power Consumed By Computers
		Amps	kVA	kW
One	276	365	101	98
Two	120	1590	190	184
Three	275	532	146	142
		Total	437	424

Summary Table of Measured Data Converted To Output Data

Computer Room	Average Volts	Total Amps	Apparent Power Consumed By Computers	Real Power Consumed By Computers
		Amps	kVA	kW
One	120	841	101	98
Two	120	1590	190	184
Three	120	1218	146	142
		Total	437	424

Appendix B: Frequently Used Terms

AHU: air handling unit

auxiliary equipment: Mechanical and electrical equipment used to support the computer equipment.

building power density: Total power drawn by the building (in watts) divided by the total floor area of the building (in square feet).

cfm: cubic feet per minute

chiller: Mechanical equipment used to make chilled water for use in cooling a building.

co-location: Refers to the act placing computer equipment owned by one company in a data center owned by a second company. In co-location facilities, the data center owner does not own the computer equipment.

computer equipment: Includes equipment such as routers, servers, hubs, switches, disks, firewalls and other information technology equipment.

computer room: Refers to the rooms in the data center hosting facilities that contain the rentable space. This term includes all of the area in this room including aisles, racks, and areas within the room that contain mechanical equipment.

computer room power density: Power drawn by the computer equipment (in watts) divided by the computer room floor area (in square feet).

computers: I use this term interchangeably with computer equipment. Includes equipment such as routers, servers, hubs, switches, disks, firewalls and other information technology equipment.

corporate data center: A data center owned and operated for an individual company (or in some cases for an individual organization or institution). Compare to hosting data center.

CRAC units: computer room air conditioning units

data center: A facility that is used to house the computer equipment to support the Internet or telecommunications system.

gross area: Refers to the total building floor area.

hosting facility or hosting data center: A data center that rents either physical or virtual (i.e., computer memory) space to its customers. A facility that "hosts" computer equipment or computer services. Compare to corporate data center.

HP: horsepower

HVAC: heating ventilation and air conditioning systems

IT: information technology

LBNL: Lawrence Berkeley National Laboratory

managed hosting facility or managed data center: A data center where the owner of the data center owns the computer equipment and rents the computer memory, function, and related services.

net area: Refers to the computer room floor area.

plug load: Electrical equipment such as lights, clocks, electric pencil sharpeners, etc. that are plugged into the electrical outlets.

PDU: power distribution unit; alternatively a power management module or PMM

total computer room power density: Power drawn by the computer equipment and all of the supporting equipment such as PDUs, UPSs, HVAC and lights (in watts) divided by the computer room floor area (in square feet).

U: The standard designation for the height of the computer equipment; 1U = 1.75 inches.

UPS: uninterruptible power supply

References

ASHRAE/ATCE, information from an anonymous HVAC engineer, 13 March 2001.

Ahlberg, Erik, "Electricity Utilities Fear Drought at Server Farms," *Dow Jones*, 30 January 2001

Angel, Jonathan, "Energy Consumption and the New Economy," *Network Magazine*, 1 January 2001.

Basler and Hofman, "Energieverbrauch von Netzwerkkomponenten (English Version)," Bundesamt für Energiewirtschaft Forschungsprogramm Elektrizität, 26 November 1997.

Bors, Douglas, "Data Centers Pose Serious Threat to Energy Supply," *bizjournal.com*, 6 October 2000.

Brown, Richard and Jonathan Koomey, "Analysis of California Electricity End-Use (draft)," Lawrence Berkeley National Laboratory, May 2001.

California Energy Commission website, "Silicon Valley Electricity Consumption," http://38.144.192.166/electricity/silicon_valley_consumption.html, viewed 29 March 2001.

Callsen, Thomas P, "The Art of Estimating Loads," Data Center Issue 2000.04, August 2000.

Cook, John, "Internet data gain is a major power drain on local utilities," *Seattle Post Intelligencer*, 5 Sept 2000.

"Council Oks restrictions on downtown 'telcom hotels,'" *Sacramento Business Journal*, 31 August 2000.

Energy Information Administration, *Annual Energy Outlook 2001*, Department of Energy Report#:DOE/EIA-0383(2001), 22 December 2000. Available at http://www.eia.doe.gov/oiaf/aeo/.

Energy Solutions and Supersymmetry, Data Center Market Research Study, Presentation for PG&E Customer Energy Management, 19 October 2000.

Feeder, Barnaby, "Digital Economy's Demand for Steady Power Strains Utilities," *New York Times*, 2 July 2000.

Globix, www.globix.com, viewed 15 March 2001.

Hall, Mark, "Net Blamed as Crisis Roils California," *Computer World*, 15 January 2001.
California Energy Commission website, "Silicon Valley Electricity Consumption," www.energy.ca.gov/silicon_valley_consumption.html, viewed 29 March 2001.

HostPro website, www.hostpro.com, viewed 12 April 2001.

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.

Kawamoto, Kaoru and Jon Koomey, Bruce Nordman, Mary Ann Piette, and Richard E. Brown, "Electricity Used by Office Equipment and Network Equipment in the U.S.: Detailed Report and Appendices," LBNL Publication 45917, February 2001. Available at http://enduse.lbl.gov/Info/LBNL-45917b.pdf.

Lazarus, David, "Net Complex a Dilemma for San Jose," *San Francisco Chronicle*, 22 March 2001. This is a compound of 10 buildings. Square footage from http://www.ci.san-jose.ca.us/planning/sjplan/eir/USDataport/US-Dataport-Text.htm.

Mahedy, Stephen and Dan Cummins and Danny Joe. "Internet Data Centers: If Built…Will They Come," Salomon Smith Barney Report, 3 August 2000. (Salomon Smith Barney Report)

McCarthy, Mike, "Ordinance would limit downtown 'telecom hotels,'" *Sacramento Business Journal*, 21 July 2000.

Mills, Mark and Peter Huber, "Dig more coal—the PCs are coming," *Forbes Magazine*, 31 May 1999.

Nordman, Bruce, "Electricity requirements for LBNL's Networking Hardware," memo to Jon Koomey at LBNL, LBNL Publication 835, 9 December 1999. Available at http://enduse.lbl.gov/Projects/InfoTech.html (Publication #3).

Nordman, Bruce, LBNL memo, 5 December 2000.

Peyton, Carrie, "Data servers crave power: High-tech electricity needs amplify crisis," *The Sacramento Bee*, 26 November 2000.

PG&E website, www.pge.com/003_save_energy/003b_bus/003blcl_program_info.shtml, viewed 12 April 2001.

Qualters, Sheri, "Energy costs surge for new projects," *Boston Business Journal*, January 29, 2001.

Reiterman, Tim, "San Franciscans Protest as 'Server Farms' Sprout," *Los Angeles Times*, 26 March 2001.

Richman, Eric E. and Carol C. Jones, and JoAnne Lindsley, "An Empirical Data Based Method for Development of Lighting Energy Standards," Journal of the Illuminating Engineering Society, Summer 1999.

Salomon Smith Barney Report. See Mahedy et al.

Schmelling, Sarah, "The Frugal Data Centers," *the Net Economy*, 2 April 2001.

Stein, Jay, "More Computing Power, Less Electrical Power," in *ET Currents* Number 7, E Source, March 2001.

The Uptime Institute, "Heat Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment," White Paper issued by the Uptime Institute, 2000.

Wilkins, Christopher and M.H. Hosni, "Heat Gain From Office Equipment," *ASHRAE Journal*, June 2000.

Yankee Group, "Executive Summary of The U.S. Collocation Market: High-Tech Real Estate Heats Up," 2000, www.yankeegroup.com/ viewed 28 March 2001.