Improving Data Center (Server Room) Energy Efficiency at Caltech

Daniel Alvarez
SURF Project 2008
August 19, 2008

1 Abstract

Caltech is at a unique turning with regards to energy usage and production. Not only is our energy consumption reaching critical loads, but we must also strive to obtain our needs from sources that are environmental-friendly and thus reduce our carbon footprint. Caltech has three major electricity conduits that each provide a maximum of 10MW. Given that we are already peaking at about 22MW and rising, this provides us with zero redundancy and a failure in any three of the conduits will result in a cascading power failure. My project is one of many steps that Caltech has taken to become more energy conscientious, specifically through improving the efficiency of Caltech's numerous server rooms. Server rooms are oftentimes over 40 times more energy intensive than a typical corporate office building. Due to the fact that our centers were built in a period where energy sustainability was the least of our concerns, the first step in my project was to benchmark our rooms through the accepted metric called the **PUE = (Total Facility Power)/(IT Equipment Power)**. Data was collected straight from our rooms using ultrasonic flow meters, amp meters, thermal guns and more to determine power consumption from IT equipment, cooling equipment, and distribution losses. The results yielded highly efficient, and unexpected, PUE values of around 1.4. Despite this there are still many improvements that can be made through hot/cold aisle isolation, energy-efficient hardware, server consolidation, proper airflow, and effective room architecture. My data supports that these improvement can in fact save money through energy savings.

2 Executive Summary

Caltech has many challenges in the future with respect to energy consumption. While needing to maintain our place at the cutting-edge of scientific research, we still need to address the imposing issue of our energy consumption reaching critical capacity. I have attempted to confront this problem via a study of the energy efficiency of four of Caltech's server rooms. At the beginning of the summer, we had no idea how efficient our rooms were, but instead had the general pretense that our rooms were relatively old and thus inefficient. My work has proven that, contrary to the beliefs of many, our rooms are in fact quite efficient through national metrics. I have found that these results are due much more in part to our highly efficient central plant chillers rather than an energy efficient architecture in our rooms and have thus been able to isolate where, and how, improvements need to be made.

3 Caltech's Energy Problem

Being at Caltech gives one the privilege, among others, to enjoy essentially unlimited amounts of energy free of charge. The natural supply-demand equilibrium of this situation is most likely the root of our current energy problems. Caltech has three major electricity conduits that can each provide a maximum of 10MW. Given that we are already peaking at about 22MW, this provides us with zero redundancy and a failure in any three of the conduits will result in a cascading power failure. This is especially true in times like this summer where our natural gas-powered cogeneration plant was not functioning due to mechanical problems. Compared to 2007 when we generated 77% of our own power, so far in 2008 we have generated only 60% of our own power. With the new campus buildings ready to come on-line soon, our energy usage is very unlikely to decrease unless we start making changes now. While Caltech's energy-feeder system can be increased in size, and indeed the feasibility of such an upgrade is being investigated, the basis on my SURF project and others is to determine whether the multi-million dollar investment can be avoided by making our current expenditure go further through increased efficiency.

I chose to specialize in improving the efficiency of Caltech's numerous server rooms. These server rooms, occupying several locations on campus, are responsible mostly for performing high-performance computing (HPC), while still powering other important things like our campus-wide IMSS networks. At the beginning of the summer it was estimated that all of Caltech's server rooms collectively consumed 2MW of power, totaling over 16% of Caltech's total power consumption per year. It was also assumed that our rooms had significant room for improvement due to their older architecture and hardware. I will later attempt to prove that this is only partially true. Despite the improvements in efficiency that need to be made to Caltech's energy usage, they must of course take a lower priority to Caltech's fundamental mission statement which states: "The mission of the California Institute of Technology is to expand human knowledge and benefit society through research integrated with education". Along these lines, a HPC Task Force Report dated on Oct. 16, 2007 indicated that in order to remain on the cutting edge of scientific research and analysis an additional 2.5MW of computing power would be required by the year 2013. This would also require further power associated with cooling, power distribution, and lighting. I believe my research has opened doors as to how we can meet these demands in the future.

4 The Energy Efficiency Problem in Data Centers

So what is a data center or server room anyway? These compacts rooms, ranging in size from less than 200 sq ft (Tier I server room/closet) to greater than 5000 sq ft (Tier IV enterprise-class data centers), are characterized by rows of racks, each holding multiple data-manipulating server blades. To put things into perspective, data centers, and their smaller server room counterparts, are oftentimes over 40 times more energy intensive than a typical corporate office building. While data centers are often located underground or in large abandoned warehouses, they keep the information technology (IT) world spinning, powering online videos to corporate e-mail.

A recent article from the New York Times illustrates the recent uproar in Silicon Valley and beyond over the newfound demand for mechanical engineers due to a lack of power in data centers. With the simultaneous rise in computing power and energy costs, the energy bill, is finally beginning to show their weight—and it is catching corporations around the world by surprise. Between 2000 and 2007 the number of data technology servers has doubled, and current trends indicate that it will double again by 2011. In the United States alone, they consumed about 61 billion kWhr of electricity in 2006, costing approximately $5 billion. Furthermore, given the fact that the United States obtains the majority of its power from burning coal, this expenditure has been estimated to consume 415 million tons of coal, and thus responsible for the release of 864 million metric tons of the greenhouse gas CO_2.

In response to the unforeseen growth in energy costs, many of the world's top IT companies including Microsoft, Sun, APC, Intel, and HP have united to form a non-profit global consortium called the Green Grid. The Green Grid's mission statement states that the consortium is, "dedicated to developing and promoting energy efficiency for data centers and information service delivery". Among the achievements of the Green Grid is, at long last, the creation of a universal metric determining the energy efficiency of a data center. This metric is known as the Power Usage Effectiveness (PUE), and it is defined as follows (See Exhibit A):

PUE = (Total Facility Power)/(IT Equipment Power)

According to the EPA's *Report to Congress on Server and Data Center Energy Efficiency* submitted on August 2, 2007, only 40% of a typical facility's power actually powers the IT equipment while approximately 30% will be responsible for cooling and the remaining 30% will be lost from associated power distribution (See Exhibit B). In terms of PUE, the conclusion is that the average data center has a PUE value of 2.5. For every 1 kW of computing power in an average center, 2.5 kW of power will be consumed in total. At Caltech, this means that hopes of achieving an additional 2.5MW of computing power by 2013 could result in an additional 6.25MW burden on our campus energy feeders.

5 The Benchmarking Problem in Data Centers

Let us go back a decade, when the IT explosion began to get out of control. Computing power had only one motto: the faster and the better. People would pay over a thousand dollars for the fastest processors available and the IT industry was in full swing, frequently reaching new orders of magnitude in processing speed. With all the new processing power, storage came to great demand and also made remarkable technological improvements. The only metric that was truly never considered was how much power each computer was draining from the wall. The last decade has seen a paradigm shift in the industry as computing power and storage, true to the merits of a free market, eventually dropped significantly in price and the growing costs of energy have become the determining factor.

A suitable analogy for the development of the server blade comes from the incandescent light bulb, which has actually changed little from Edison's original design despite the fact that 90% of the energy consumed is given off as heat. This is because the design is very *effective*, though not at all *efficient*. The same applies to most IT servers today. The relation between IT load and power consumption is incredibly inelastic, meaning that even essentially idle servers (5-10% IT load) are still using approximately 60% of their full capacity power. The compromise between efficacy and efficiency is finally shifting as more IT managers are realizing they can either not obtain or not afford the energy associated with expected IT growth trends.

As aforementioned, the benchmarking problem is rooted from the fact that the data center stage was set back when power consumptions was last, if at all, on the list of priorities. For that reason, the primary objective of my work was to analyze and determine the PUE of four of our largest server rooms, essentially "benchmark" our current state of affairs with respect to energy efficiency.

6 Caltech's Server Rooms

While there is a cornucopia of literature available on-line describing Tier III and IV data centers like those which Google and Microsoft would operate, there was in fact very little dealing with the Tier I server rooms located at Caltech. Naturally the dynamics differ significantly as a 1% efficiency increase in one of Google's massive data centers is an investment worth making, while the initial capital required to obtain the same improvement here is doubtfully worth the investment. The small size of our server rooms means that every one of them differs significantly both from each other and from the typical data center set-up (see Exhibit C). These differences include whether they have an uninterruptible power supply (UPS), power distribution units (PDUs), or none of the above. Caltech has many server rooms and closets ranging from a few blade servers stacked up in a professor's office to the "Beowolf" facility located beneath S. Mudd, which is a largest on campus. For this reason, I chose to investigate four of Caltech's largest rooms: S. Mudd 0064, CACR, IPAC 104, and Thomas 012. I also chose these rooms due to their greatly differing architecture, which is easily the largest factor affecting how efficiently the room is powered and cooled. Placing the racks in a "mirror image" alignment, for example, so that the heat-emitting backs of servers face each other creates hot/cold aisles in data centers that prevents cold air from being re-cooled. Thomas and S. Mudd 0064 are more modern and have well-implemented (completely implemented in Thomas's case) hot/cold aisle isolation while the aged architecture of CACR and IPAC yielded very poor hot/cold aisle isolation.

The results of my research support aforementioned conclusion (See Exhibit D). However, the results ranging from PUE values of 1.4 to 1.6 seemingly suggest that there is little difference between the rooms and that they are all very efficient. I must confess that I was initially very puzzled as to how such remarkable efficiencies could be obtained. Eventually I realized that the low PUE values were due to a factor outside the scope of the rooms: very efficient chillers cooling our campus's water supply. So efficient, in fact, that Mr. Reza Ohadi, the director of campus operations, flew to Washington D.C. to receive an award for operating one of the most efficient plants in the country. Due to the fact that it only takes us .7kW per ton of cooling, you may notice that our chilled water expenditures are significantly less than those expected of a typical data center, ranging from just 9-15%. The two examples at the bottom of exhibit D from *Self Benchmarking Guide for Data Center Energy Performance* performed by the Laurence Berkeley National Labs display a typical inefficient and highly efficient data center. In both cases the power expenditures associated with chilled water far outweigh the fan power expenditure associated with the Computer Room Air-Conditioners (CRACs) located inside the rooms. In our rooms, the chilled water and fan power are far more balanced. Also, because of the fact that our rooms are smaller and can afford down-time (unlike, for instance, a banking or credit card center) we get away with no back-up generators, few memory units, and, in cases like Thomas, the lack of both a UPS and PDUs. This eliminates the natural distribution losses associated with changing the voltage and/or redistributing it. Finally, it must also be noted that the PUE is less sensitive to change than just a straight percentage (also known as the DCE) with Thomas's PUE value of 1.4 representing 72% of the total power going to IT equipment and IPAC's PUE value of 1.6 going representing 65% of the total power going to IT equipment, a significant difference.

6 The Road to Efficiency

Unlike original estimates at the beginning of summer indicating that our server rooms consume about 2MW total, my research has shown that the four rooms I investigated in fact consume only 815kW, allowing a safe assumption of 1MW including the rooms I did not take into account. This would indicate that our server rooms consume just over 8% of Caltech's yearly power consumption. Despite impressively low PUE values in our centers when compared to on-line literature on the subject, all of are rooms continue to have room for improvement.

On a national scale, the EPA report to congress gave several wedge analyses associated with certain improvements that can be made to data centers around the country (See Exhibit E). Among the most important recommendations that can be made to any data center is server consolidation and utilization of efficient hardware. Due to limitations in our size, Caltech will never host the world's largest computational resources. The San Diego Supercomputing Center (SDSC) and NASA Center for Computational Sciences (NCCS), among others, are where our scientists go when such top-end computing is required. Our size limits our ability to consolidate our centers into larger and thus more efficient ones. However, even within our current server rooms existing servers can and should be consolidated with more efficient hardware. One case study I examined found that a company is saving $1 million per year in energy costs after they consolidated 368 servers running at an average 17% IT load to 5 servers running at an average 70% IT load. Steps have been taken in the right direction at Caltech. In the "Beowolf" room in S. Mudd, for instance, the number of racks was halved after existing hardware was replaced by state-of-the-art quad-core servers. These servers effectively triple the performance of a single server while utilizing virtually the same amount of energy through sharing internal components. Of course, it might be necessary to implement certain incentives to encourage our faculty to buy more efficient, but also more expensive, hardware over time.

Besides the Beowolf room modifications in S. Mudd, many other changes are taking place around campus to accommodate the computing needs of the future. The basements of Jorgenson and Keith Spalding are being rehabilitated and turned into server rooms and CACR and IPAC 104 are both soon going to be remodeled to implement completely isolated hot/cold aisles. At the beginning of the summer there was no accurate measure of the efficiencies in these rooms. Temperature and humidity sensors are now being installed in the rooms, along with flow and temperature sensors on the chilled water pipes, to keep real-time estimates of every room's PUE. The analysis that I have performed will give clear before/after comparisons in these rooms once they are refurbished. In Thomas, space has been specifically left aside to easily accommodate new sets of racks within the completely isolated hot/cold aisle design. While the isolated hot/cold aisle solution will undoubtedly help through preventing cooled air to be re-cooled, airflow problems will begin to arise as soon as new racks get closer to the CRAC units. The next step will be to balance the air flow using "In row" CRAC units which blow air directly from hot aisle to cold aisle (See Exhibit F). Balanced airflow will ensure that no "hot spots" arise in the server room and thus allow the overall temperature to be raised, reducing the workload of the CRAC units. The server blades are guaranteed to function properly at 79ºF, but rooms currently lacking hot/cold aisle isolation are blowing air as cold as 60 ºF to ensure that all hot-spots in the room are at most 79ºF. Ensuring balanced airflow with In-Row CRAC Units and isolated hot/cold aisles eliminates hot spots and allows the overall temperature to be raised to 75 ºF. Once this is accomplished, air economizers can also be purchased to blow in "free" outside air at night which more often than not drops below 75 ºF.

The final wedge on the EPA analysis, and thus also called the "state-of-the art" solution is liquid cooling racks which run chilled water directly through custom-made computer racks. Because water has over a thousand times the heat capacity of air, these racks can efficiently cool computer racks internally. We have recently purchased 4 rows of liquid-cooled computer racks for the server room that is to be constructed below the new Annenberg IST building. I was personally disillusioned, however, by the fact that they are unnecessarily (in my opinion) installing a supplemental Liebert CRAC unit for, according to the project manager, general room and lighting cooling. Despite the fact that server rooms should run dark as often as possible, I believe that even purchasing more efficient lighting (LEDs or T-5 fluorescent lights) would be a far better investment than an expensive supplemental unit.

While liquid cooling racks and In-Row computing power can be increased while electrical power is decreased. Despite the associated initial capital investments, the energy savings easily add up hundred of thousands of dollars within years (See Exhibit G). Such changes will keep Caltech at the cutting-edge of research, save money, and ensure the future generation a cleaner world.

6 References

Dunlap, Kevin, and Neil Rasmussen. The Advantages of Row and Rack-Oriented Cooling for Data Centers. White Paper #130. APC.

Hannaford, Peter Top Ten Cooling Solutions to Support High-Density Server Deployment. White Paper #42. APC

Kutz, Andrew. Improving Data Center Energy Efficiency: A Holistic Approach Burton Group. V1.0

Mispagel, David California Institute of Technology HPC Feasibility Study. Caltech. July 2008

Petrino, Micheal, and McHugh, Daniel J. A Comprehensive Report of Findings & Recommendations for California Institute of Technology. PTS Data Center Solutions. November 27, 2007

Report to Congress on Server and Data Center Energy Efficiency. Environmental Protection Agency. August 2, 2007

Self Benchmarking Guide For Data Center Energy Performance Ernest Orlando Lawrence Berkeley National Laboratory. V1.0

7 Appendix

Exhibit A: PUE

PUE: Power Usage Effectiveness
DCE: Data Center Efficiency

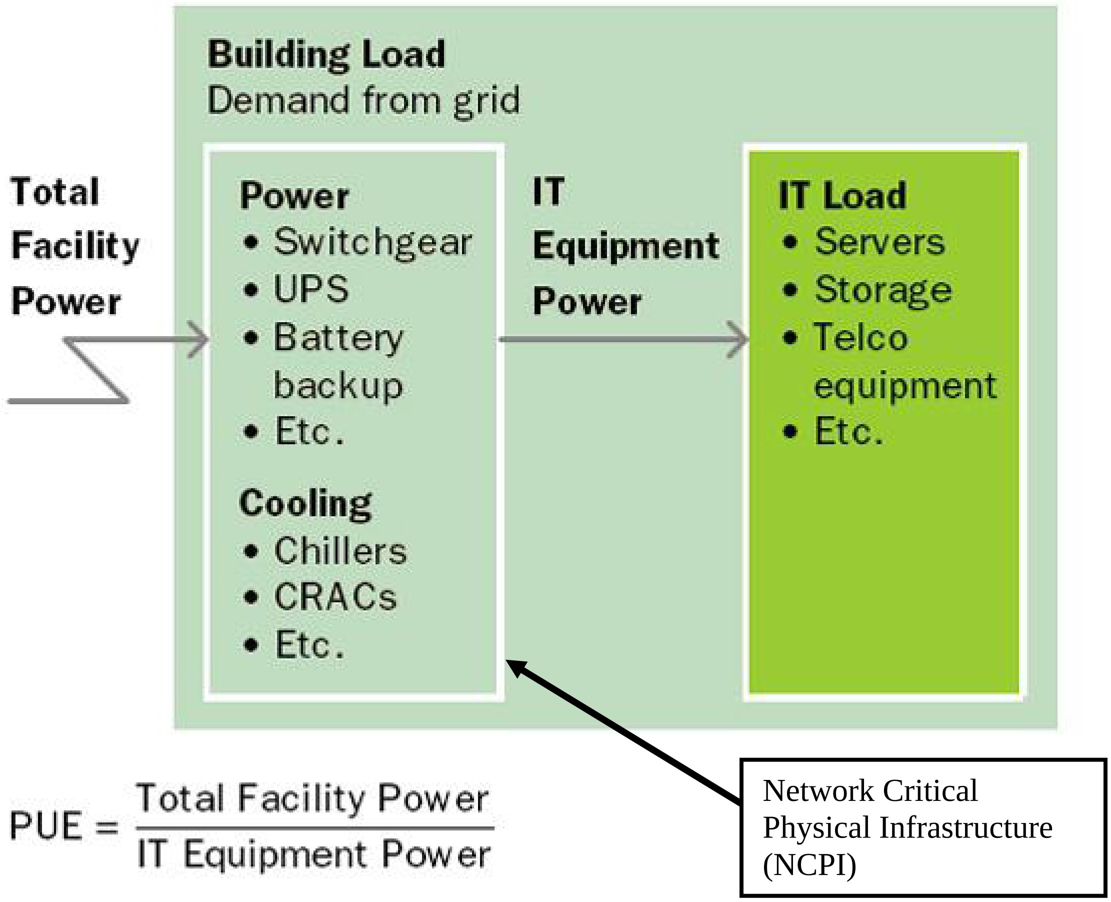

$$PUE = \frac{\text{Total Facility Power}}{\text{IT Equipment Power}}$$

$$DCE = \frac{1}{PUE} = \frac{\text{IT Equipment Power}}{\text{Total Facility Power}}$$

Exhibit B: Average Hierarchal Data Center Power Distribution

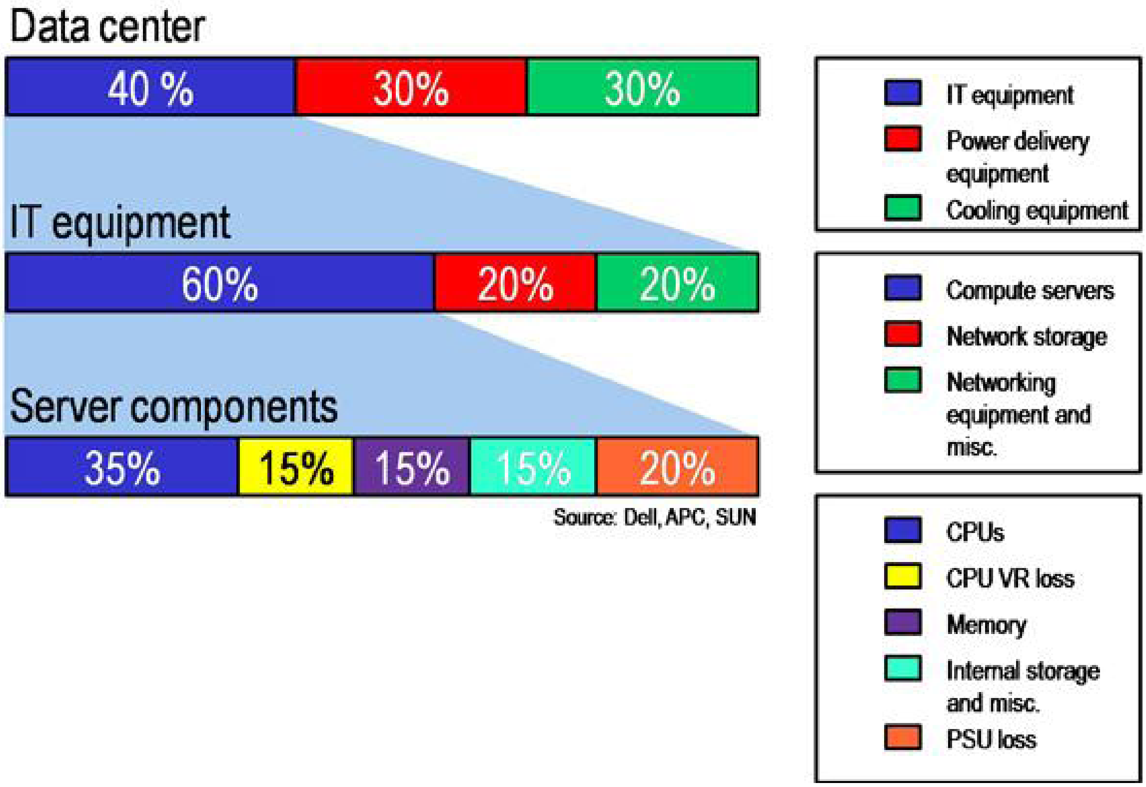

Exhibit C: Typical Electrical Flow in a Data Center

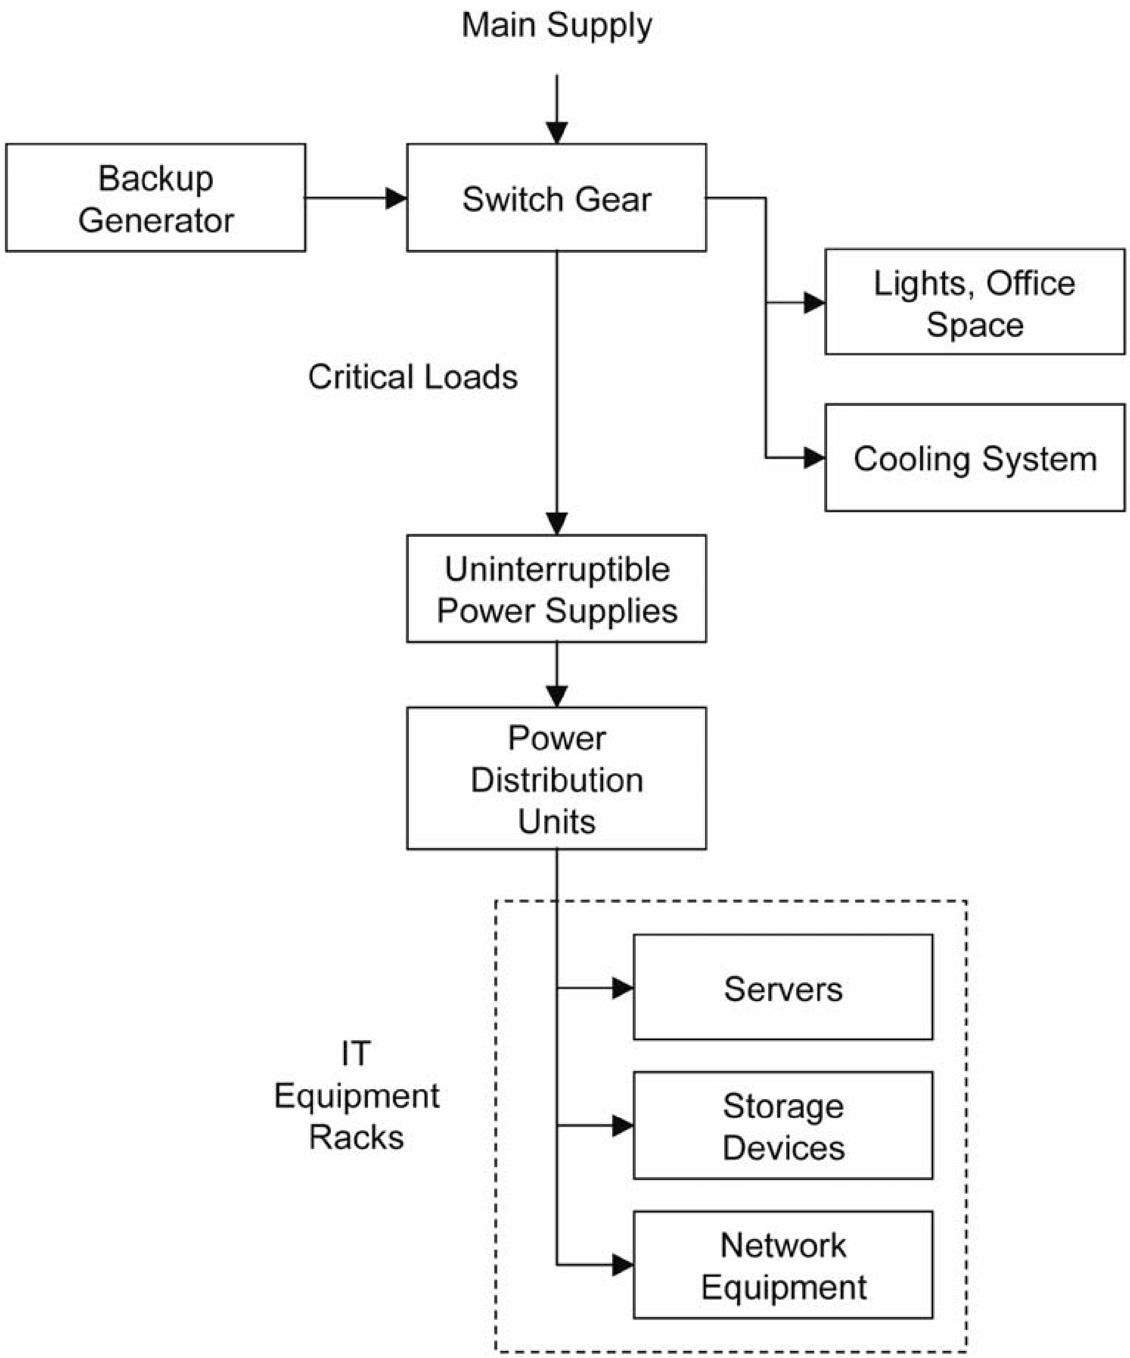

Exhibit D: Power Distributions

IPAC 104 Power Distribution

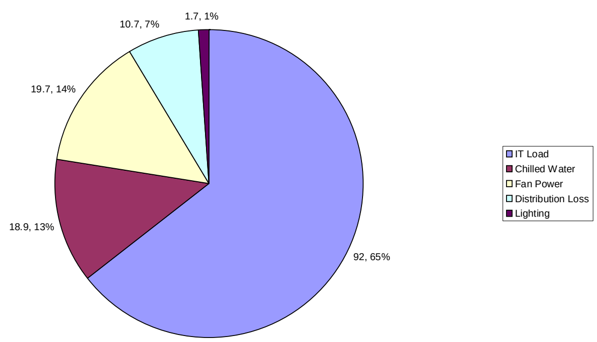

PUE = 1.6

CACR Power Distribution

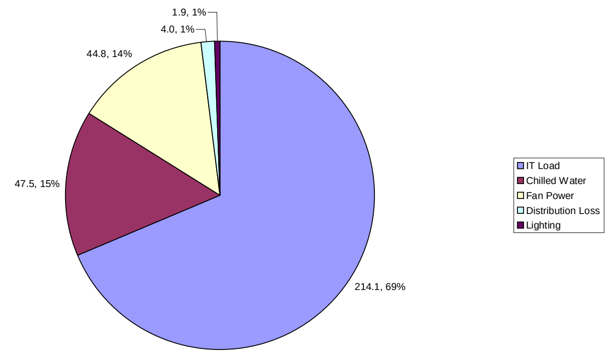

PUE = 1.5

S. Mudd 0064 Power Distribution

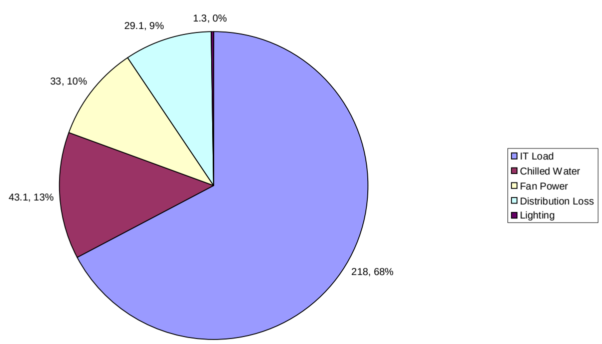

PUE = 1.5

Thomas 012 Power Distribution

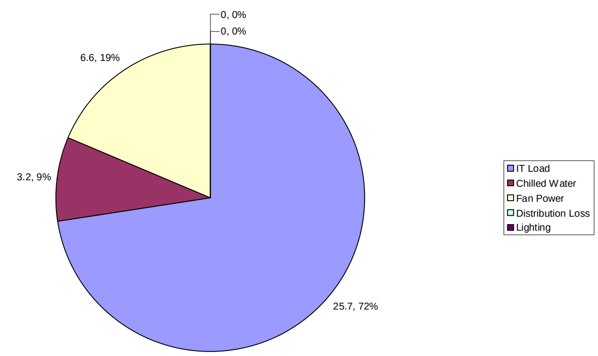

PUE = 1.4

Inefficient Room

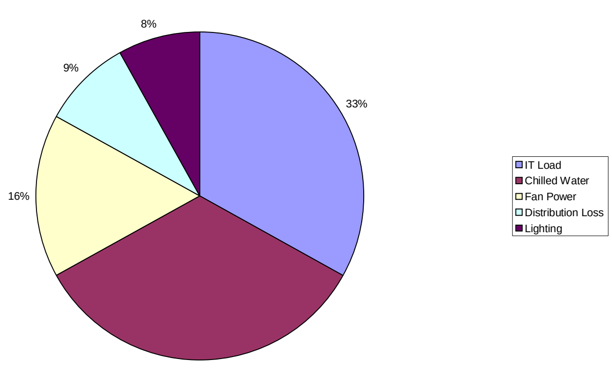

PUE = 3.0

Efficient Room

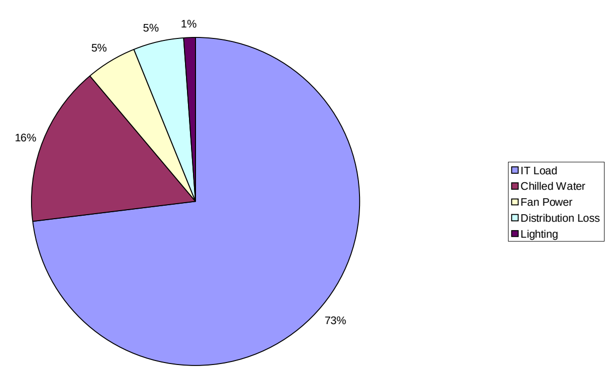

PUE = 1.4

Exhibit E: Official ENERGY STAR (EPA) Recommendations to Congress and their Potential Savings

Data Center Subsystem		
Scenario	**IT Equipment**	**Site Infrastructure (Power and Cooling)**
Improved Operation	• Continue current trends for server consolidation • Elliminate unused servers • Adopt "energy efficient" servers • Enable power management on all servers • Assume modest decline in energy use of enterprise storage equipment	30% improvement in infrastructure energy efficiency from improved airflow management
Best practice	All measures in "Improved operation" scenario, plus: • Consolidate servers to moderate extent • Aggressively adopt "energy-efficient" servers • Assume moderate storage consolidation	Up to 70% improvement in infrastructure energy efficiency from all measures in "Improved operation" scenario, plus: • improved transformers and uninterruptible power supplies • improved efficiency chillers, fans, and pumps • free cooling
State-of-the-art	All measures in "Best practice" scenario, plus: • Aggressively consolidate servers • Aggressively consolidate storage • Enable power management at data center level of applications, servers, and equipment for networking and storage	Up to 80% improvement in infrastructure energy efficiency, due to all measures in "Best practice" scenario, plus: • direct liquid cooling • combined heat and power

Scenario	Electricity Consumption Savings (billion kWhr)	Electricity cost savings ($ billion 2005)	Carbon dioxide emissions avoided (MMTCO$_2$)
Improved operation	23	1.6	15
Best practice	60	4.1	38
State-of-the-art	74	5.1	47

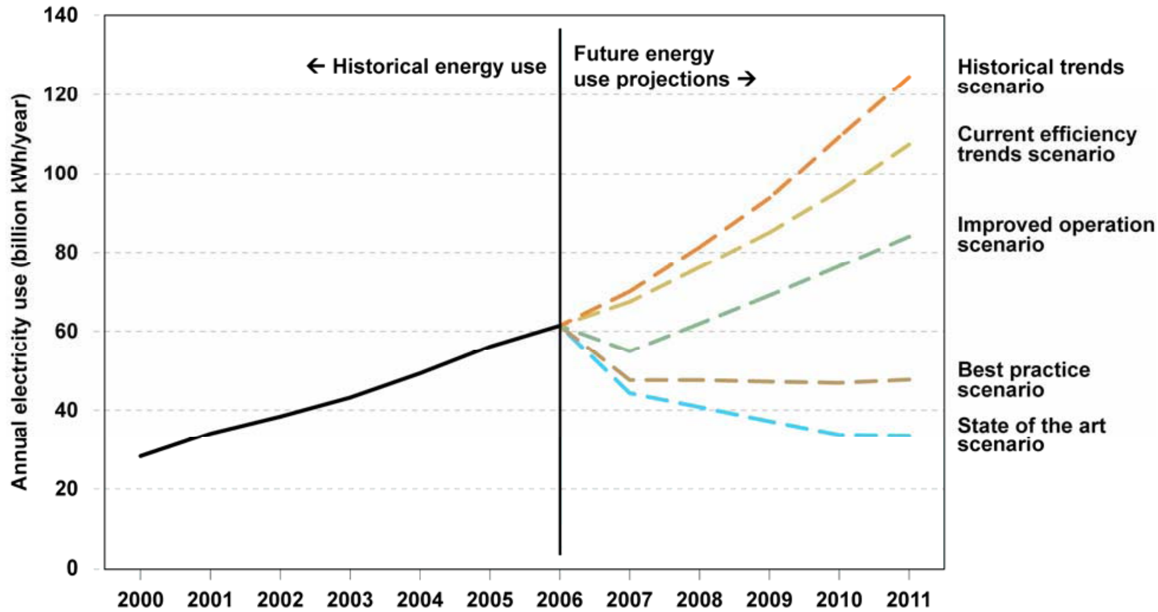

140
120
100
80
60
40
20
0
Annual electricity use (billion kWh/year)
← Historical energy use
Future energy use projections →
Historical trends scenario
Current efficiency trends scenario
Improved operation scenario
Best practice scenario
State of the art scenario
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011

Exhibit F: Two Alternative Cooling Methods: Liquid Cooling and Targeted Cooling

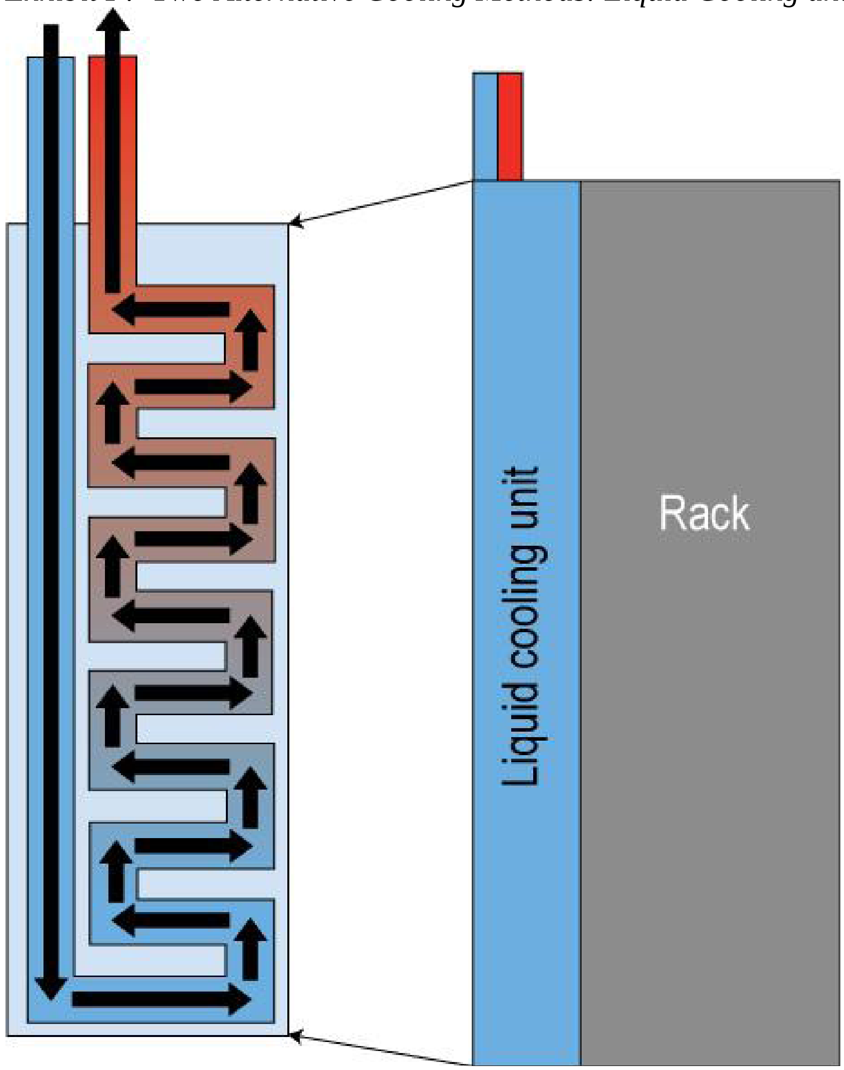

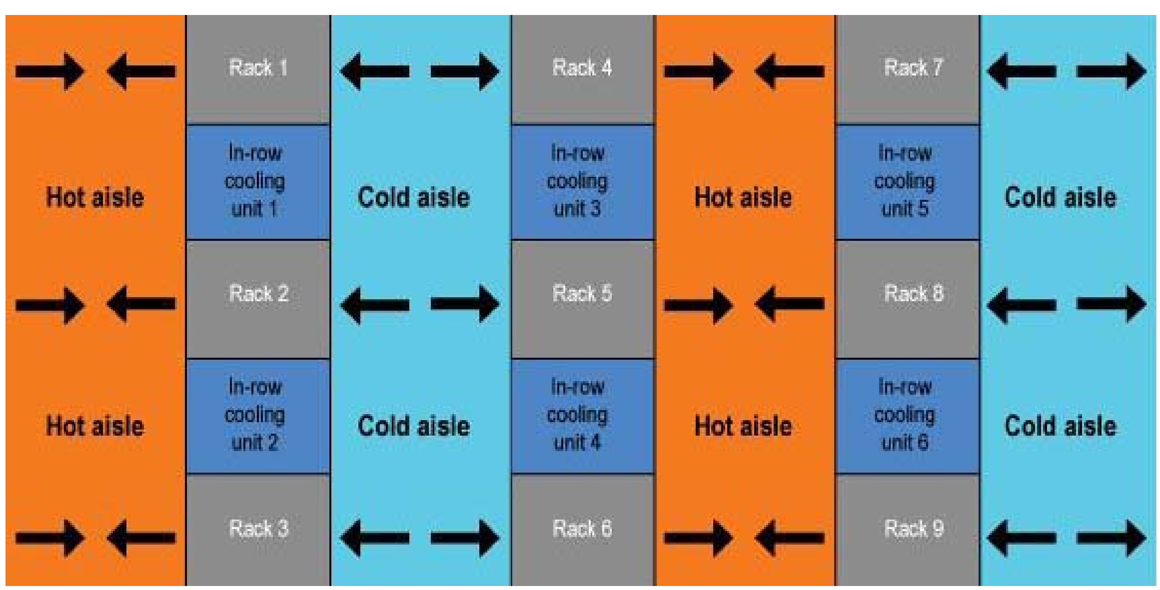

Exhibit G: Annual Savings at PWP cost of $.15/kWh

Load (kW)	PUE: 1.4	$	PUE: 1.6	$	PUE: 1.8	$	2	$
250	350	459900	400	525600	450	591300	500	657000
1000	1400	1839600	1600	2102400	1800	2365200	2000	2628000
2500	3500	4599000	4000	5256000	4500	5913000	5000	6570000