DC Power for Improved Data Center Efficiency

March 2008

My Ton

Ecos Consulting

Brian Fortenbery

EPRI

William Tschudi

Lawrence Berkeley National Laboratory

The work described in this report was performed as part of Lawrence Berkeley National Laboratory's (LBNL) "High-Performance High-Tech Buildings" project focused on improving energy efficiency in Laboratory, Cleanroom, and Data Center facilities. This project was sponsored by the California Energy Commission's Public Interest Energy Research (PIER) Industrial program with technical leadership by LBNL. Ecos Consulting and EPRI Solutions performed the research reported herein collaboratively under subcontract to LBNL and with input from industry advisors and manufacturers.

Legal Notice

Table of Contents

1. Executive Summary

1.1 Introduction

The mission of data centers in use today calls for them to be much more energy intensive than other office buildings. Recent reports prepared by this project team for the California Energy Commission estimated that the US annual energy consumption by the equipment in these centers can use up to 14.6 TWh, with supporting equipment such as uninterruptible power supplies (UPS) accounting for up to 7.1 TWh. The data center market represents an important component of the California economy as well as a considerable and growing source of electrical demand. Preliminary investigations confirmed that research with the objective of improving the efficiency of data center was warranted.

Typical data center power delivery designs use AC (alternating current) power, distributed to the facility at 600V AC or 480V AC. This AC power is then stepped down to 208V AC or 120V AC for distribution to racks for use by servers and other information technology (IT) equipment. An UPS and energy storage system, such as batteries or flywheels is used to isolate equipment from power interruptions or other disturbances. This set up generally involves converting incoming AC power to DC (direct current) for energy storage. The DC power is then converted back to AC for the facility distribution grid and routed to power distribution units (PDUs) for distribution to equipment in racks.

Inside the servers and other IT equipment such as storage or networking units, power supplies convert AC (at 208/120V AC) to DC voltage needed for the digital electronics. Power supplies usually provide power factor correction as well as load isolation from the incoming power line for these sensitive electronic components. Thus, there can be up to six or more power conversion stages between facility power entry and the microprocessor or other data processing circuits.

The power losses due to the use of inefficient power conversion devices from both outside and within equipment result in a large loss of useful electrical power, as well as directly increasing the energy required to remove the heat produced. Thus, for every watt of power utilized to process data, about 0.9W is required to support power conversion. In addition, about 0.6 to 1 watt will be required to cool the power conversion equipment.

This report details a demonstration of alternative approaches for delivering power to computational and network equipment in a data center using DC, and comparing its efficiency to the more traditional approach using AC.

1.2 Project Objectives

The objectives of this demonstration project are to develop and demonstrate a power delivery system that does not contain as many power conversion stages using existing equipment and vendors where possible.

This project implemented a power delivery system that distributes DC to the server racks. The system used a single rectification stage, thereby removing the conventional UPS, transformer, and the rectifier in the server's first stage power supply. A standard AC distribution system is installed next to this DC system, server loads were connected and programmed to run identical routines. For this identical amount of computing work, the input power for whole system was measured and compared.

The following goals were identified for this demonstration project:

1. Show that DC-powered server(s) exists in the same form factor or can be built.
2. Show that DC-powered server(s) provides the same level of functionality.
3. Measure and document any efficiency gains from the elimination of multiple conversion steps in the delivery of DC power to the server hardware.
4. Identify areas requiring further development or follow up investigations.

1.3 Project Implementation

As part of the previous server power supplies and UPS work, the project team identified a number of industry stakeholders and held discussions with them regarding participation in a possible demonstration of an integrated DC power delivery system.

In early 2006, a suitable location for the demonstration project was found in Newark, CA. The project participants also helped to define the three configurations for the demonstration project. They are:

- **AC Reference Configuration**: This configuration is needed to simulate current data center typical set-up, delivering 208/120V AC input to AC-powered servers, and to be used as a reference to compare conversion efficiency.
- **Facility-Level DC Configuration**: This configuration is needed as the proof of concept – the ability to deliver high-voltage DC throughout the facility. This configuration handles the DC conversion/distribution at the building/data center level, converting 480V AC to 380V DC and delivering this directly to the DC-powered server units in the rack.
- **Rack-Level DC Configuration**: This configuration is needed to provide a possible migration option for AC data centers operators wishing to use DC equipment without facility-wide DC power distribution. This configuration accomplishes DC conversion/distribution at the rack level, using a rectifier unit to convert 208/120V AC at the rack, and delivering 380V DC to DC-powered servers.

In addition, a number of conditions were agreed upon by the participants and the project team on implementation, including:

- **Testing and Measurements**: The group agreed on test points and metrics, with emphasis on measuring the efficiency of the configurations – there would be no direct comparison of server equipment performance.
- **DC Input Voltage**: Due to compatibility with existing equipment and devices, the group settled on 380 V DC for the high voltage DC input.

Project participants contributed time, equipment, in kind services or input towards project implementation. The demonstration set up as defined was completed by mid June 2006, and demonstrated to the press and other industry representatives on June 21. A series of "open houses" were held through out the months of July and August. Over 200 visitors attended these demonstrations.

The interest generated by the project helped to bring about other manufacturers of DC equipment, notably manufacturers of communication and support equipment using 48 V DC. These manufacturers provided additional computing and telecommunication equipment for the use of the demonstration project. The presence of this equipment provided another reminder that DC power distribution is not new, and has been safely and effectively used in telecommunication and data networks. In addition, it also showed that DC and AC power delivery system can co-locate in a data center facility.

1.4 Project Results

Our results indicate that the DC approach does provide an increase in conversion efficiency. We were fortunate enough to have access to two AC distribution systems as well as two DC conversion/distribution systems, and the efficiency ratios were determined for both sets.

Table ES1

System Efficiency	UPS Efficiency	Transformer Efficiency	PS Efficiency	System Efficiency
AC System A: Measured Efficiency	90%	98%	90%	79%
AC System B: Measured Efficiency	90%	98%	90%	79%
DC System A: Measured Efficiency	94%	100%	92%	87%

Energy Consumption	Compute Load (kWh)	Input Load (kWh)	Efficiency Gain
AC System A: Measured Consumption	23.3	26.0	
AC System B: Measured Consumption	23.3	25.9	
DC System A: Measured Consumption	22.7	24.1	
% Energy Consumption Improvement vs. AC System A			**7.3%**
% Energy Consumption Improvement vs. AC System B			**7.0%**

It can be seen (Table ES1) that there is about a 7% decrease in input energy using the first DC system compared to the "best in class" AC systems. With the second DC system, the values are slightly lower, but still about 5% improved over the AC systems (Table ES2).

Table ES2

System Efficiency	UPS Efficiency	Transformer Efficiency	PS Efficiency	System Efficiency
AC System A: Measured Efficiency	90%	98%	90%	79%
AC System B: Measured Efficiency	90%	98%	90%	79%
DC System B: Measured Efficiency	92%	100%	92%	85%

Energy Consumption	Compute Load (kWh)	Input Load (kWh)	Efficiency Gain
AC System A: Measured Consumption	23.3	26.0	
AC System B: Measured Consumption	23.3	25.9	
DC System B: Measured Consumption	22.7	24.6	
% Energy Consumption Improvement vs. AC System A			**5.0%**
% Energy Consumption Improvement vs. AC System B			**4.7%**

It should be noted that both of the AC distribution system used represent the best on the market with regard to efficiency. Both of the AC UPSs are high efficiency units, and the efficiencies of the power supplies in the AC servers – at 90%, are much higher than units currently found in today's data centers. By comparison, a typical AC system in today's data center would have a UPS that was about 85% efficient, and power supplies around 73% efficient.[1] The estimated improvement of the DC system over these "typical" systems is shown in Table ES3 below.

Table ES3

System Efficiency	UPS Efficiency	Transformer Efficiency	PS Efficiency	System Efficiency
AC Typical Distribution Efficiency	85%	98%	73%	61%
DC Distribution Efficiency	92%	100%	92%	85%

Energy Consumption	Compute Load (W)	Input Load (W)	Efficiency Gain
Typical AC Distribution Efficiency	10,000	16,445	
DC Distribution Option (Optimized)	10,000	11,815	
% Energy Consumption Improvement vs. Typical AC Distribution			**28.2%**

In this case, an improvement of over 28% is possible in an average data center. This means the DC distribution system, as demonstrated, will have the potential of using 28% less energy than the typical AC system found in today's data centers. Since data center HVAC loads are typically about the same as the IT load, this means that a 28% improvement in distribution and conversion also means a 28% overall facility level efficiency improvement.

It should be noted that the magnitude of the DC efficiency gain is highly dependent on the AC reference system and AC/DC power supply that it is being compared to. However, exposing the industry to "best in class" systems may also be useful in bringing attention to the need for improving the average efficiency of data centers.

1.5 Conclusions and Recommendations

This demonstration project was able to coordinate the participation of 21 organizations, their equipment, and/or in kind contribution, worked with other organizations' input throughout the implementation process, and assembled equipment and services worth over a million dollars in value. We were also able to conclusively demonstrate to the data center industry (via the 200+ open house attendees and the media coverage) that DC delivery systems are viable, can be 20% or more efficient than current AC delivery systems, be more reliable, and potentially cost less in the long run.

Overall, the project succeeded in meeting the objectives that were set out at the beginning. In particular:

Availability of DC Equipment: The demonstration project showed that DC-powered servers exist in the same form factor as AC servers or can be built and operated from existing components with minimal effort. DC servers currently exist (in the 48V DC form

[1] These are typical numbers that were found in our evaluation of the servers and UPSs markets.

factors), but 380V DC servers and storage equipment could be built and operated from existing components. Further, the demonstration project gave visibility to the DC power conversion and distribution equipment, highlighting two commercially available rectification systems, as well as UL-listed buss bars for DC applications.

DC Functionality: The project also showed that DC-powered servers can provide the same level of functionality and computing performance when compared to similarly configured and operating server containing AC power supplies. The demonstration equipment included storage units as well as DC network equipment that can use a variety of DC voltages.

Demonstrated Gains in Efficiency: The project demonstrated clear efficiency gains from the elimination of multiple conversion steps in the delivery of DC power to the server hardware. Results were measured and documented from two sets of DC delivery systems, and compared to two sets of AC delivery systems. In both cases, the DC delivery system showed a minimum of 5% to 7% efficiency gains without significant optimization over two AC distribution systems that are "best in class" and much more efficient than most systems found in today's data centers. These measured efficiency gains did not include additional gains from a reduction in cooling loads, which can have the potential for additional savings. Raising awareness of the AC - UPS system efficiency will have a benefit even if the DC solution is not embraced.

Follow-Up Investigations: A number of areas for follow up investigations were identified that will help generate industry discussions, and provide useful leverage points to move the industry forward in the direction of DC distribution. These include:

- **Grounding, Protection and Overloading Prevention**: A number of grounding, protection and overload prevention practices for DC data centers are discussed in this report for industry considerations.
- **Reliability**: Anecdotal data shows that DC-powered data centers have the potential to be more reliable than AC-powered data centers. However, data does not exist for DC voltages higher than 48V DC. Leadership will need to come from the industry in adoption or additional testing for the industry to move forward on this area.
- **Costs**: While cost data exists, it has not yet been compiled in a way that direct comparisons can be made for the two distribution systems (or their TCO). The demonstration has generated significant interests from data center designers and system integrators, and further discussion on this area can lead to at least a first-order estimate of DC distribution costs. A related area is the costs of DC components vs. AC components. Currently, AC components may enjoy better economies of scale, but wide-spread DC power adoption may change this equation.
- **Integration with Other Sources**: The use of DC in data center can also simplify the integration of alternative energy sources, such as solar and other forms of renewable energy, as well as fuel cells and distributed generation, which are all DC-based.
- **Other Issues**: Of significant concerns is the lack of industry knowledge of the advantages of DC distribution, as well as misconceptions about DC power. Additional education and outreach efforts will be required if the energy savings potentials of DC powered data centers are to be realized.

Going forward, there remain many barriers to the adoption of DC power distribution, and need additional follow on work. In particular, a number of barriers have been identified during the course of this projects, they include:

- **Increasing Awareness of DC Distribution**: The industry's current knowledge of available options for efficiency and DC distribution is quite low. Further, there is no single, trusted source of information, or an entity dedicated to the promotion of data center energy-efficiency and DC distribution (other than the CEC/PIER efforts). With the current industry support and interest, a "DC Power" association of some sort will help to focus interest and help to elevate awareness among the data center market.
- **Creating a Market for DC**: A number of market barriers still need to be addressed in a consistent, unified manner. Coordinating utilities and other efforts, at least in California, will go far towards getting DC approaches to take hold in the data center market. There is utility interest in establishing a baseline of performance and cost, which can then help to address at least the early adoption barrier of cost. Other efforts are still needed, and strategies to address market transformation used by the conservation movement can be directly applied here. In addition, the US Congress has recognized the potential for energy savings with HR-5646, so that coordination with DOE and EPA is needed to ensure no duplication of efforts.
- **Develop standards to accelerate adoption**: Agreement on distribution voltages, electrical connectors, grounding, DC power quality, and other issues will be important to enable the market to adopt DC distribution on a large scale. The PIER program should facilitate these efforts by bringing together the appropriate industry representatives.
- **Developing Pilot Projects**: Once the Demonstration Project was completed, there are no other places where such a set up can be found. Efforts are needed to continue the Demonstration Project's role in informing the industry of the DC distribution alternative. Discussions are underway with a number of "early adopters," and the project team proposes to establish several pilot projects to:
    - Create demand for DC servers to enable certification efforts to proceed
    - Determine cost factors for DC systems - capital and operating cost
    - Evaluate and resolve any remaining barriers
    - Publicize successful systems in real data centers.

2. Background

2.1 Introduction

This report details a demonstration of alternative approaches for delivering power to computational and network equipment in a data center using DC (direct current), and comparing its efficiency to the more traditional approach using AC (alternating current). The report is divided into five main sections: The first section discusses data centers' energy requirements, power delivery systems and efficiency implications. The second section deals with the project chronology and definition of project objectives and equipment. The third section covers issues encountered during the implementation. The fourth section discusses safety and other associated issues. The fifth section provides the project results, both technical and non-technical. The last section summarizes the conclusions, recommendations, and next steps.

2.2 Project Background

There are many types of "data centers" in use today across the U.S. and around the world.[2] They range from corporate data centers in a wide range of industries: banks, telecommunication facilities, to Internet services facilities, as well as institutions such as research organizations, universities, national laboratories, and government facilities. They generally consist of a collection of servers – high powered computers typically housed in equipment racks, in tandem with networking and storage equipment to store and move large amounts of data. Along with the data processing equipment are the supporting equipment: power distribution, back-up systems, cooling, and other services to maintain data center operations.

A number of other definitions have also been used to describe "data centers" and their functions.[3] At the most basic level, data centers house computers and information technology (IT) equipment to provide functions such as data information storage, data processing, and information dissemination. With the boom associated with the Internet in the late '90's came new names for data centers including "server farms", "collocation facilities" and "telecommunication hotels."

Regardless of their make-up or affiliation, most data centers' mission calls for them to be much more energy intensive than other office buildings. This is due to the high power requirements of not only the computing equipment, but also for the infrastructure needed to support the computing equipment, and their constant availability, or "up time." In fact,

---

[2] The terms "data centers" and "datacenters" are used interchangeably by the industry. This report uses "data centers." For the purpose of this report, these are dedicated facilities that comprise more than 1,000 square feet of computational space and support equipment.

[3] See, for example, http://en.wikipedia.org/wiki/Data_center, or ACEEE and CECS 2001; Aebischer et al. 2002b; Blount et al. 2001; Brown et al. 2001; Callsen 2000; Elliot 2001; Gruener 2000; Mitchell-Jackson 2001; Planet-TECH 2002; Robertson and Romm 2002.

large data centers more closely resemble industrial facilities in their energy density than typical commercial buildings.

Figure 1. Data Center

Source: UC Berkeley

The data center market has a significant presence in California, and represents an important component of the California economy as well as a considerable and growing source of electrical demand on the state's electric utilities. Preliminary investigations and utility case studies confirmed that research with the objective of reducing the large, continuous electrical loads in data centers was warranted. Because of this, California utilities and the California Energy Commission (CEC) became interested in learning more about the high-tech sector and the data center market.

As part of the California Energy Commission's PIER (Public Interest Energy Research) initiative on improving energy efficiency for high-tech buildings, Lawrence Berkeley National Laboratory (LBNL) was tasked with preparing R&D roadmaps for data centers as well as for cleanrooms and laboratories, focusing on reducing their energy use.[4] Using these roadmaps, sponsors of public sector R&D programs are able to prioritize their investments in the research, development, and demonstration of innovative energy efficient technologies and best practices for the next decade.

2.3 Overview of Data Center Power Use

From the beginning of commercial computing, data centers have been important to industries, institutions, and governmental agencies. However, it was the Internet and the rise of "mission-critical" computing facilities that brought energy consumption in data centers to the forefront. In 1999, a widely cited article by Mark Mills and Peter Huber posited that the "Internet" consumed about 8% of US electricity in 1998, and that this

[4] Tschudi, W.F., Xu, D. Sartor, and D. Stein. 2003. "High Performance Data Centers: A Research Roadmap". LBNL-53483. Available from http://hightech.lbl.gov/library.html#Publications.

would likely grow significantly to account for 30 to 50% of U.S. electricity consumption by 2010 or 2020. They cited the dramatic growth in Internet-related equipment installed base since 1995, particularly servers and computer network equipment needed for the Internet.[5]

Mills and Huber's work generated significant interests by the US Congress, the US Department of Energy (US DOE) and prompted additional investigations by other researchers. Lawrence Berkeley National Laboratory and Arthur D. Little, Inc., who conducted more detailed studies (Koomey et al., 1999; Koomey, 2000; and Roth, Goldstein, and Kleinman, 2002) conclude that this initial estimate was highly overstated, and that computer, office, and network ICT equipment potentially accounted for only about 3 percent of U.S. power consumption in 1999/2000.[6]

Common among these follow on studies is the observation that there is a large variation in energy intensity and energy efficiency of key systems in the various facilities under examination. It is safe to say that current understanding of data center energy use in the industries and institutions that rely on them is very limited. Typically, data center professionals have a thorough understanding of issues related to power quality, reliability, and availability, as these have been high priority items, but not energy consumption or efficiency.

To accurately estimate the power consumption and future needs of data centers in the US, one place to begin is to characterize the stock of existing data centers and their load intensity. These characteristics have turned out to be difficult to estimate. The market is a competitive one, and can be constantly changing. Further, there is no reliable source of market data covering all of the various types of data centers, and companies operating them tend not to publicly disclose their capacities and loading. Load intensity for data centers supporting the Internet fluctuates greatly with the rise and decline of internet-based companies. However, data center load intensity is also affected by the trends in computing capability and energy intensity within IT equipment.

There are three main energy trends that have been identified by studies as drivers for this new awareness regarding data center energy requirements:

- **Advances in computer technologies**: faster chip technology was creating higher heat density in smaller and smaller geometries. The simultaneous compaction and increase in electrical power caused concern over the ability to cool future generations of IT equipment.

---

[5] Mills, M and P. Huber, "Dig More Coal—the PCs are Coming," Forbes, May 1999.

[6] Roth, K; F. Goldstein; J. Kleinman "Energy Consumption by Commercial Office and Telecommunications Equipment in the U.S." Presentation at "Impact of ICT on the Energy System" 2002.

- **Significant increases in electrical power requests from utilities by facilities**: meeting growing requests by computing facilities could require major changes in current electrical utility generation and distribution infrastructure. [7]

- **Predictions of large increases in electrical demands for future facilities**: IT professionals, data center operators, and facility designers have been predicting even larger requirements for future computing facilities.[8]

Ecos Consulting and EPRI Solutions were tasked with investigating and refining existing estimates of the overall energy requirements by this sector as a follow up to the above studies in this area. In particular, Roth, Goldstein, and Kleinman in 2002 arrived at an estimated Annual Energy Consumption (AEC) for IT equipment using a bottom-up analysis of the stock, usage patterns and power draw (by mode) of IT equipment. They concluded in their analysis that for 2000, the estimated AEC was about 10.1 TWh for US data centers, as shown in Table 1, below.[9] Their methodology is as follows:

- Generate a list of equipment types and collect existing data from literature.
- Develop a preliminary estimate of national energy consumption for each equipment type.
- Select 5 to 10 equipment types for further evaluation, based upon preliminary calculations and perceived growth in future energy consumption. Ideally, the selected equipment types should represent 66% to 75% of all energy consumed by office and telecommunications equipment in the commercial buildings sector.
- Briefly describe the 5 to 10 equipment types selected, with the intent to provide insight into how each equipment type uses energy and function in a commercial office environment, including: physical description, functions performed, and commercialization history.
- Develop refined bottom-up estimates of national energy consumption of each selected equipment type, for Y2000 and projections for Y2005 and Y2010.
- Compare the results with the results of other studies.
- Qualitatively discuss possible indirect impact of commercial office and telecommunications equipment upon energy consumption, e.g. e-commerce, building heating and cooling loads, etc.
- Publish the findings in a report, including feedback from government and industry experts.

It is important to note that the Roth's methodology established here for estimating the AEC, and also the follow up work by Ecos and EPRI only account for the direct energy consumption of the computing equipment (servers/computational equipment and UPSs)

---

[7] Tschudi, W.F., P. Sreedharan, T. Xu, D. Coup, and P. Roggensack. 2003. "Data Centers and Energy Use – Let's Look at the Data." ACEEE Conference Paper.

[8] Markoff, J. and S. Hansel, "Hiding in Plain Sight, Google Seeks More Power" *The New York Times*, June 14, 2006.

[9] 1 TWh = 1,000,000,000,000 Wh, or 1 billion kWh.

and not the additional energy required to support or cool the facilities.[10] Benchmarking performed by LBNL, while illustrating a wide range of performance, shows that it is not uncommon that infrastructure load requirements are of similar magnitude to the IT equipment load.

Table 1. Estimated 2000 Annual Energy Consumption for the IT Sector

Equipment Type	US AEC in TWh	% of Total AEC
Computers (between $25k and $349k)	24	22%
Computers (<$25k, i.e. PCs)	18	16%
Monitors	18	16%
Computers (>$349k)	10	9%
Laser Printers	9.3	8%
Copiers	8.3	7%
Server Computers	7.7	7%
Telecomunications Network Equipment	7.1	6%
Computer Network Equipment	6	5%
Uninterruptible Power Supplies (UPSs)	2.5	2%
Total Equipment	**110.9**	**100%**

Source: Roth et al, 2002.

Ecos' and EPRI Solution's follow up investigations of the AEC primarily focused on the various equipment and components in use in data centers and their energy consumption, with the main focus on server and equipment power supplies as well as uninterruptible power supplies.[11] For server power supplies, the goals for the project included an analysis of power supply efficiencies, which required the development of an accepted test protocol for server power supplies, lab and field testing of a broad range of server power supplies and documenting the results. Using the efficiency findings and market data, we arrived at estimates of overall energy consumption of servers in the United States, especially in the State of California (as well as the potential savings from the use of more efficient units).[12]

Similarly, for the UPS market, the investigation included an analysis of UPS topologies and a complete analysis of distribution of UPS design efficiencies, which included the development of a UPS test protocol. Using the efficiency findings and market data, we estimated the overall energy consumption of UPSs in the United States, especially for the

---

[10] Note: In this report, several terms are used generically, including "servers" as a generic term to describe the computing equipment used in data centers, including racks of servers and contents; and "UPSs" to describe any number of back-up power sources and approaches to maintaining "up time".

[11] Ecos Consulting and EPRI Solutions. 2005. "High Performance Buildings: Data Centers -- Uninterruptible Power Supplies (UPS), and "Ecos Consulting and EPRI Solutions. 2005. "High-Performance Buildings: Data Centers -- Server Power Supplies."

[12] Other objectives of the project included working with industry groups to press the case for more efficient power supplies. the wide circulation of efficiency findings to the industry through the CEC, PIER, Lawrence Berkeley National Laboratories (LBNL), and other industry and efficiency venues, such as SSI – the Server System Infrastructure group, the PSMA – the Power Sources Manufacturers Association, the American Society of Heating, Refrigeration and Air Conditioning Engineers – ASHRAE, and www.efficientpowersupplies.org.

State of California, as well as the potential savings from the use of more efficient units. Finally, the investigation also proposed efficiency levels and a labeling scheme for various types of UPS units as a way to encourage the use of more efficient UPS units. The results from the server power supplies and UPS studies are presented below.

Annual Energy Consumption of Servers (Data Processing Equipment)

Starting from the basic methodology established in the ADL/US DOE study, Ecos and EPRI Solutions constructed a revised estimate of the annual energy consumption of servers in the United States. Test results for this work showed that most server power supplies' efficiency at converting AC to DC typically peaks at loads between 50-60% and drops off dramatically at loads under 30%. The tested power supplies have efficiencies in the 70-75% range (at 50% load). Figure 2 below shows the efficiency range of the server power supplies tested.

Figure 2. Power Supply Efficiency Curves

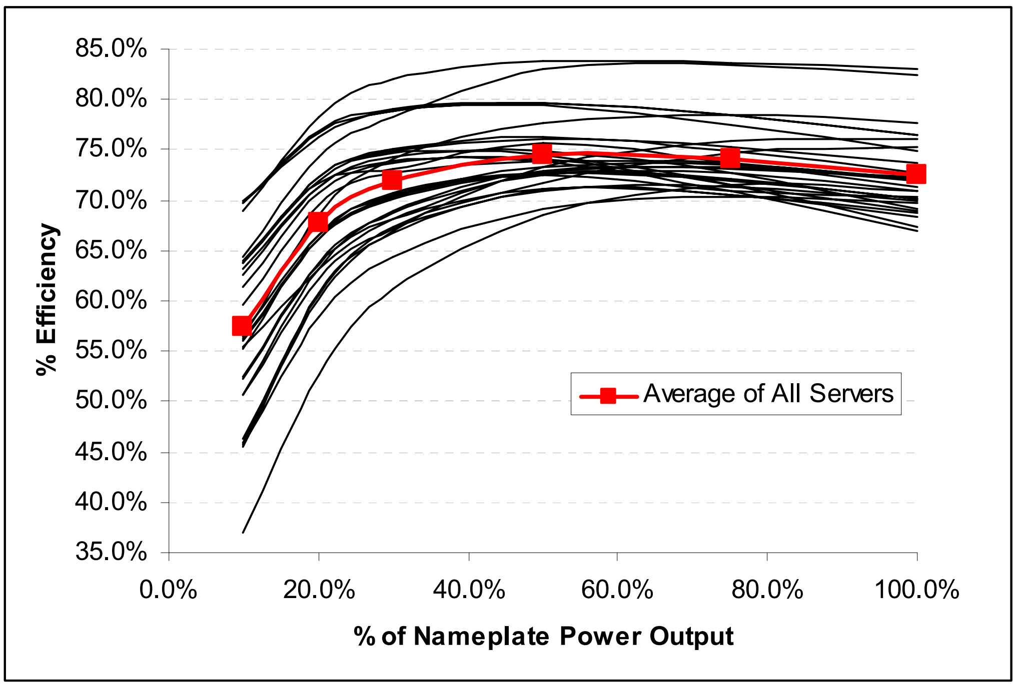

The research also showed that server power supply designs with poor efficiency are still available and can result in unnecessary power consumption and excess heat generation while in use. Further, the servers that were surveyed do not exhibit any activity-based power management. The absence of this process in the servers that were measured might help to explain the lack of any correlation between the utilization curve of the CPU and the load duration curves of the power supplies.

Figure 3 below shows our revised estimate of server AEC for 2004. The revised AEC is 14.6 TWh, which represents a 45% increase over the ADL/US DOE estimate of 10.1 TWh for 2000. This increase is attributable to the growth in both the number and average power draw of low-end servers, which comprise the majority of units. The estimate for California is in the range of 1.5 TWh to 2.2 TWh.

Figure 3. Estimated Annual Energy Consumption of Servers

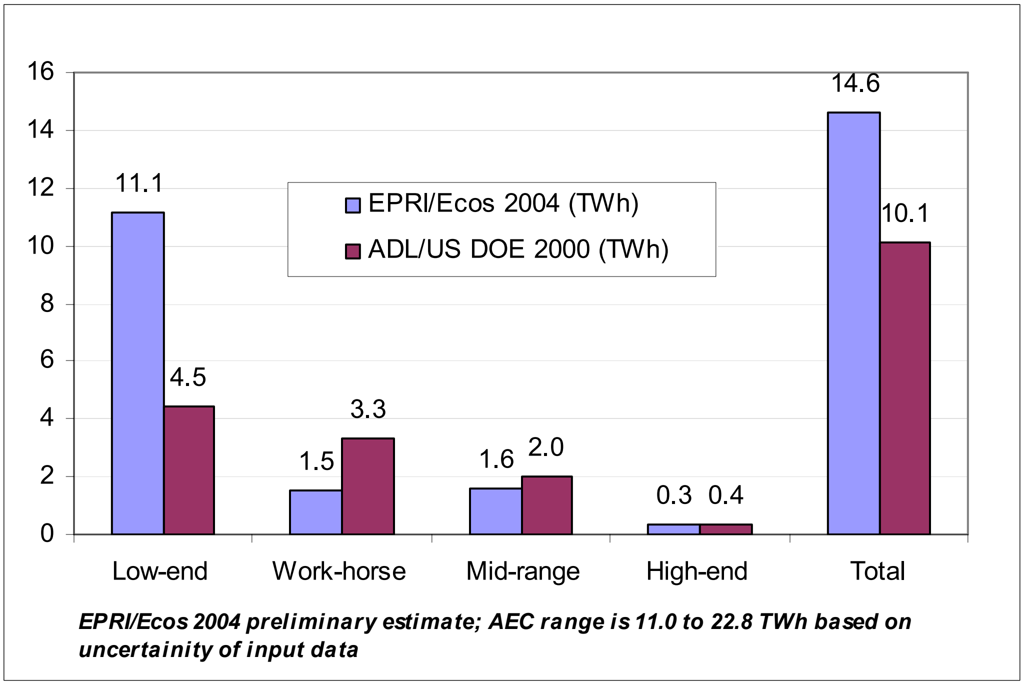

Annual Energy Consumption of UPSs

To estimate UPS power consumption, Ecos and EPRI started by sizing the UPS market and the installed stock of UPSs in the data center/IT sector. Our results indicate that there can be a wide variation in efficiency between various UPS configurations. The table below summarizes the characteristic efficiency of a number of UPS topologies at various load conditions and shows the average efficiency for all of the UPSs measured.

The table below also shows that regardless of the configuration or type, UPSs tend to be more efficient at full rather than part load which is where they typically operate due to over sizing and redundancy strategies.

Table 2. UPS Topologies and Efficiencies

UPS Topology	Efficiency at 25% Load	Efficiency at 50% Load	Efficiency at 75% Load	Efficiency at 100% Load
Delta-Conversion	93% - 94%	96% - 97%	97%	97%
Double-conversion	81% - 93%	85% - 94%	86% - 95%	86% - 95%
Line-Interactive	NA	97% - 98%	98%	98%
Standby	NA	NA	NA	NA
AVERAGE OF ALL UNITS	86%	89%	90%	90%

Note: A straight average was used in the table above.

Using efficiency data and our estimates of UPS stock, we constructed a first-order estimate of UPS energy use for the U.S. This energy use and savings estimate is limited to the data center/IT sector due to the scope of our research, in which we examined the efficiency of data centers only. Based on our estimates, UPSs in the data center/IT sector currently handle roughly 7.1 TWh (7.1 billion kWh) of electricity per year. We estimate that the State of California alone, which contains roughly 15% of the nation's data center floor space, consumes about 1 billion kWh of electricity and spends about $100 million per year in electric bills and a significant amount is due to power conversion losses. Our national energy use figures are noticeably higher than past estimates made by Arthur D. Little for the U.S. Department of Energy, which estimated UPS energy consumption in the U.S. "IT/telecom" sector at 5.8 TWh per year.

The quantification of the annual energy consumption for data centers – other than providing a plausible range and a first-order estimate – remains an elusive goal for a number of reasons. However, with the additional information gathered from the server power supplies and UPS studies, a number of observations can be made:

1. Data centers, both in the US and California, can account for a significant demand on electrical power, based on the number of servers in operation.

2. When supporting equipment are taken into account, the range of AEC for data centers become very large, even without taking their cooling needs into consideration.

3. Based on findings from our investigations into this market, there are opportunities for more efficient equipment to reduce this very large amount of energy annually needed to operate data centers.

4. Because more efficient equipment tend to reduce losses, especially conversion losses in the power delivery chain, increases in equipment efficiency can translate directly into reduction in cooling loads, leading to more energy savings. Another potential benefit is a reduction in heat-related failure.

2.4 Overview of Current Data Center Power Delivery System

In order to understand the opportunities for improving data center efficiency, it is necessary to examine the current data center power delivery system from end to end. That is, an examination of the delivery chain starting from the facility level and ending at the processor level. There are often several levels of power conversion occurring in data center facilities and within the IT equipment contained therein, as well as redundancies and other accepted practices by data center designers and builders to increase availability, or "up time". This examination will not only help to identify areas where there can be areas for improvements, but may also pinpoint issues and other opportunities for simplification. While we have identified power supplies and UPSs as two significant areas based on our research, they by no means are the only areas where possible efficiency improvements can be made.

Data Center Electrical Distribution
Current data center power delivery designs use AC power, distributed from the utility (or "at the curb") to the facility at 600V AC or 480V AC depending on the size of the facility. This AC power is then stepped down to 208V AC or 120V AC via transformers for distribution to server racks for use by the servers and other data center equipment. AC is also used in powering the ancillary support equipment, such as HVAC and lighting.[13] In the U.S., servers typically take 208V AC (or 120V AC) input. Figure 4, below illustrate this power delivery and use system.

Figure 4. Power Distribution for Data Center Equipment

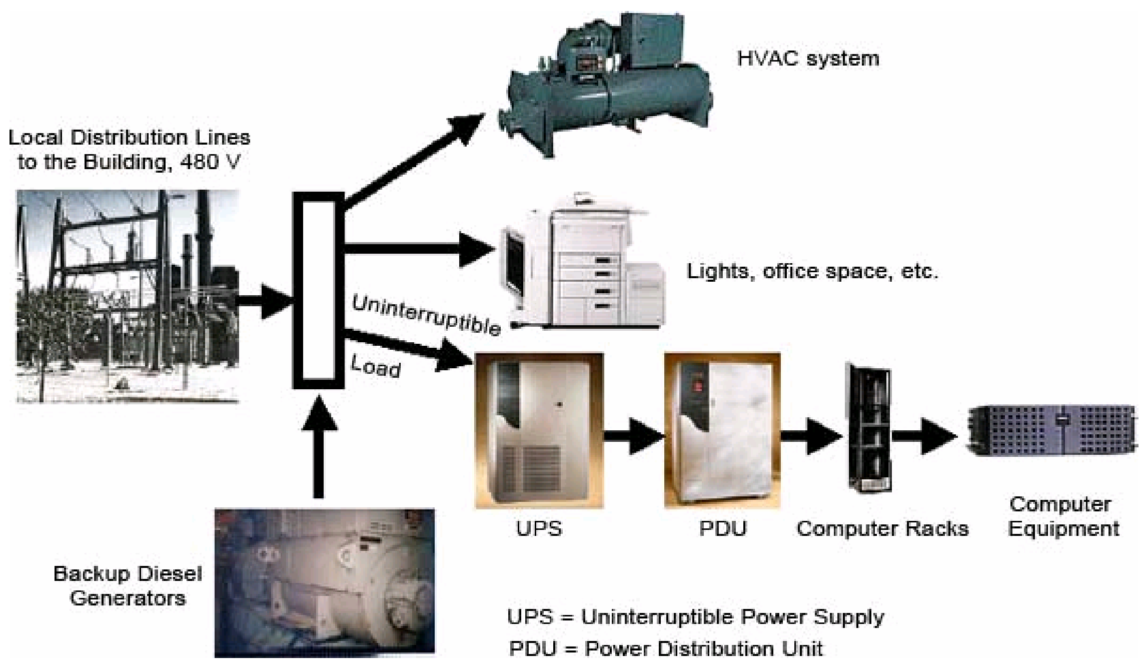

In this power chain, there is usually an uninterruptible power supply (UPS) which is coupled with and energy storage system, such as batteries or flywheels to ensure the servers and associated processing equipment are never exposed to power interruptions or other power line disturbances. Depending on the UPS configuration, the central AC UPS generally involves conversion from incoming AC power to DC for energy storage, and then reconverting from DC back to AC for distribution throughout the facility via power distribution units (PDUs).

Moreover, if the UPS requires a maintenance bypass switch, facility-level entry power must be transformed from 480V AC down to 208V AC prior to input to the UPS. The central battery is then connected as an alternate input to the DC-AC converter so that if incoming raw AC power is interrupted, the system automatically switches over to battery

[13] To insure availability, data center HVAC systems can also be on a UPS protected circuit.

power. Uninterruptible AC power is then passed through an AC power distribution grid and routed to PDUs for distribution to individual data processing apparatus such as equipment racks. The servers and other equipment in the racks receive the AC output of the UPS (converted from DC) and then convert again to DC within its power supply. Figure 5, below presents a simplified view of this power delivery chain.

Figure 5 . Simplified Illustration of Power Conversion Steps

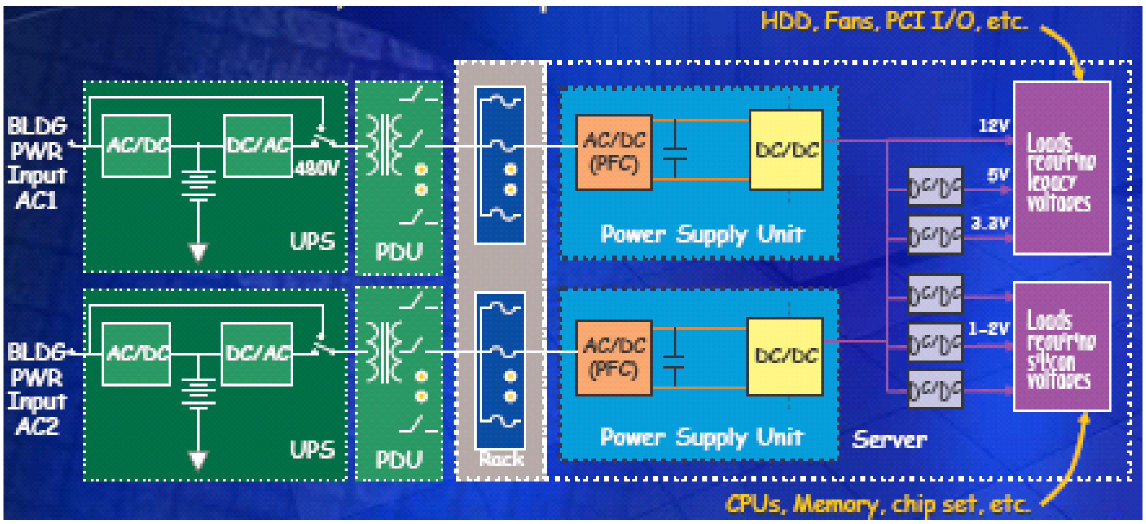

Source: Intel Corp.

The power quality and reliability requirements are expensive, but in data centers, like no other industry, these are paramount, and the extra cost is borne without question. In fact, the PDU is often connected to two parallel UPSs in some data center designs, to provide redundancy for higher availability in case one UPS fails or taken down for maintenance service or repair. In addition to the UPS redundancy, the power supplies in the servers themselves are often redundant, with two or even three power supplies in each server box ready and capable of powering the server completely in case one or more of them fails.

The typical large data center AC power distribution system is very complex and costly. Because of the reliability requirements of these facilities there are many redundancies, power conversions, paralleling controls, static transfer switches and bypass connections built into the system. It is because of these complexities that a data center power system is so much more reliable then a typical commercial building power system. It is also because of these complexities that a current day data center power system is very inefficient in both space utilization and energy efficiency.

A detailed example of the power distribution system approach (one of many) in a data center is shown below.

Figure 6. Power Delivery System Details

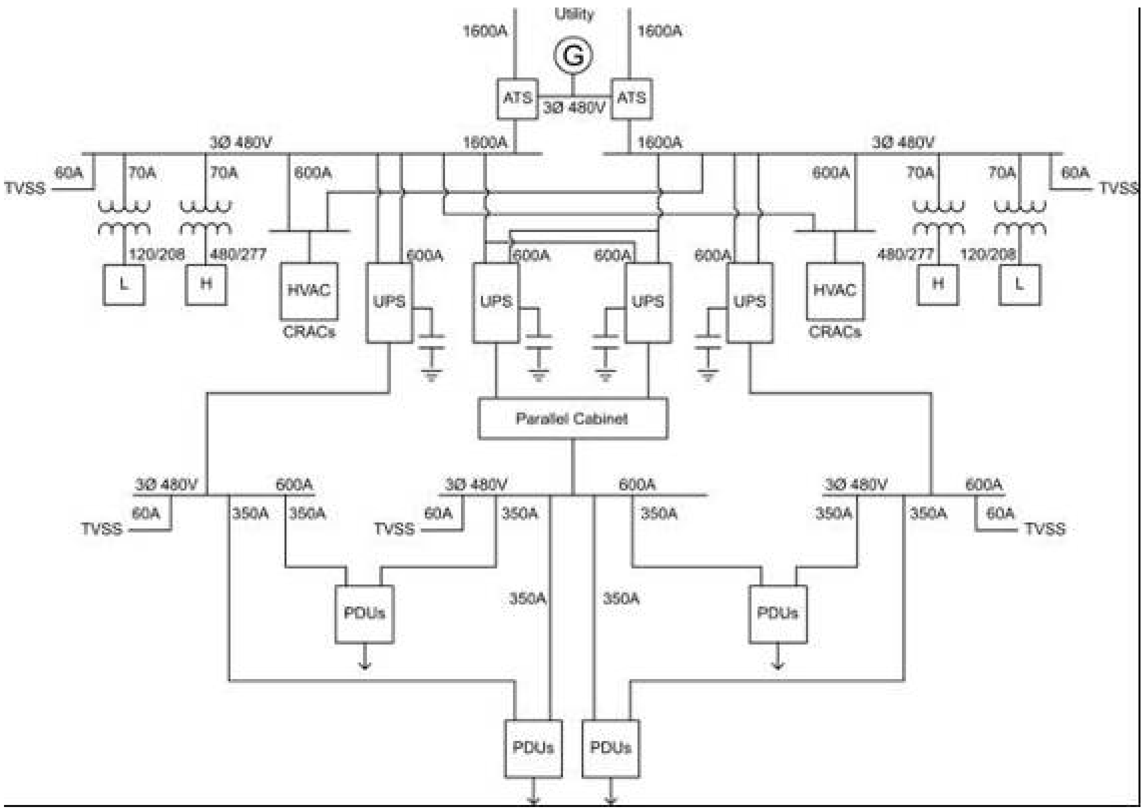

Typically, uninterruptible power is viewed at the facility level, as opposed to the equipment level. This perspective provides an easy division between the facility power equipment and data processing equipment, with each focusing on a different part of the power delivery. However, this view also makes it difficult to ascertain and optimize overall operating efficiency and total cost of ownership, since losses from each of these conversion processes directly translate into heat, adding to both the server cooling load and ultimately, the overall data center's cooling load.

Server Internal Electrical Distribution

Inside the servers and other data processing equipment such as storage units, power supplies convert AC voltage (at 208/120V AC) to DC voltage for use by the digital processing electronics. Power supplies usually provide power factor correction as well as load isolation from the incoming power line for these sensitive electronic components. This conversion generally involves at least two stages. In addition, most modern microprocessors require very low voltages at fairly high currents, such as 1.1V at 100A. The precision of the voltage required, coupled with losses incurred by distributing high power at ~1V, and the high number of different rail voltages needed dictate that voltage regulation circuitry must be located directly next to the microprocessor. In order to effectively realize this circuitry, most processors require that an intermediate DC voltage, such as 12V DC, be delivered to the processor/local regulator combination.

Thus, from the power supply, there can be up to six or more power conversion stages between facility power entry and the microprocessor.[14] Depending on processor and server loading, which can dynamically range from 30% to 100%, the efficiencies of power supplies will likely be lower at lower load levels and can significantly reduce overall system efficiency. In many cases, redundant power supplies are used to deliver this power, either on standby or in load-sharing configurations inside servers, both of which reduce individual power supply loading and efficiency.

Issues and Limitations of Current Data Center Designs

Due to the many levels of power conversion occurring in current data center facilities and within the IT equipment contained therein, most data centers experience significant electrical power losses in their facility systems' supply and distribution. This also includes losses in transformers, power line conditioners, UPS, line losses, etc.[15] Intel Corporation surveyed these losses at an actual data center, and the magnitudes of these losses for this facility are detailed in Figure 7 below.

Figure 7. Source and Magnitude of Data Center Losses

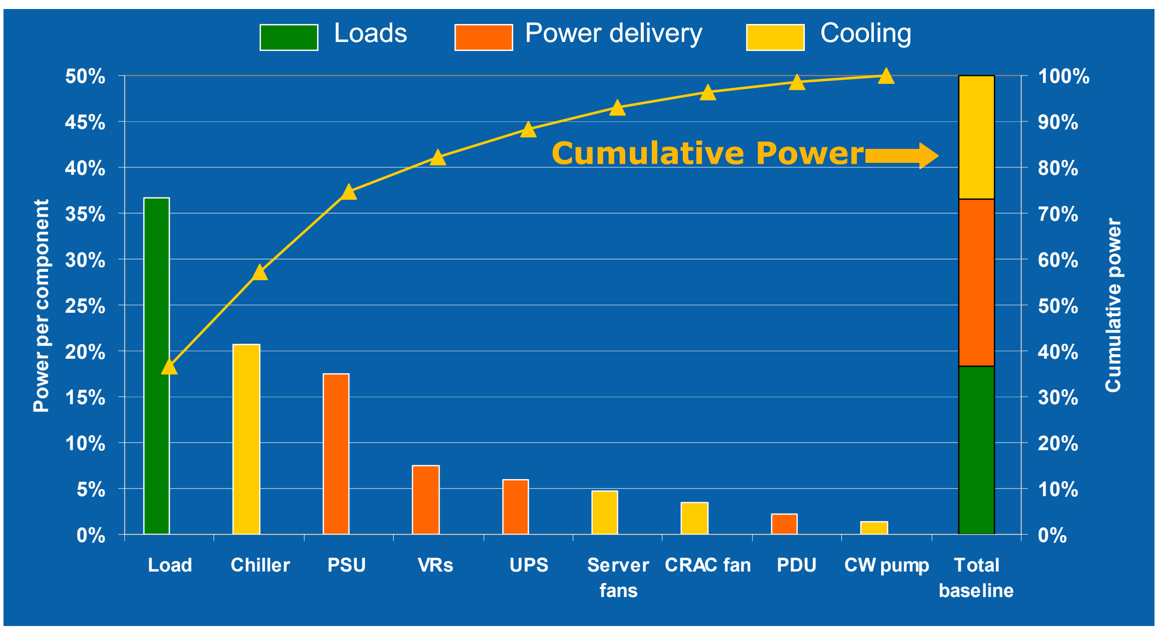

Source: Intel Corp.

The centralized battery plant utilized in facility-level UPS systems can be a compromise between what is required for successful system realization and component limitations. The actual battery run time required for acceptable system operation is often anywhere from 10 to 20 seconds to just a few minutes. The time represents the delay needed to switch to alternate utility power feeds, or to bring an auxiliary source of power on line (such as an engine generator). However, when batteries for centralized AC UPS systems

[14] Example conversion steps: 480VAC to 208VAC, 208VAC to 400VDC, 400VDC to 208VAC, 208VAC to 400VDC, 400VDC to 12VDC, 12VDC to 1.1VDC.

[15] Note that power line losses are not covered in the scope of this report.

are sized, the nature of the voltages required or battery type chosen can often result in hold up times well in excess of what is required, resulting in wasted energy storage and delivery capabilities. [16]

Inside the server units, the limiting factors on power conversion efficiency are voltage rating of semiconductors, along with their corresponding conduction losses.[17] In addition, power converter package size limitations, and the economics of power supply design and manufacturing can also play a part. All of these factors combine to ultimately limit available efficiency improvements for the power conversion processes. Power conversion efficiency for best-in-class computer grade AC-DC power supplies currently do not provide efficiency higher than about 90%. For servers, on-board non-isolated converters that provide final processor power conversion can range over 90% efficiency.

Figure 8. Conversion Losses and Their Effects on Data Center Cooling

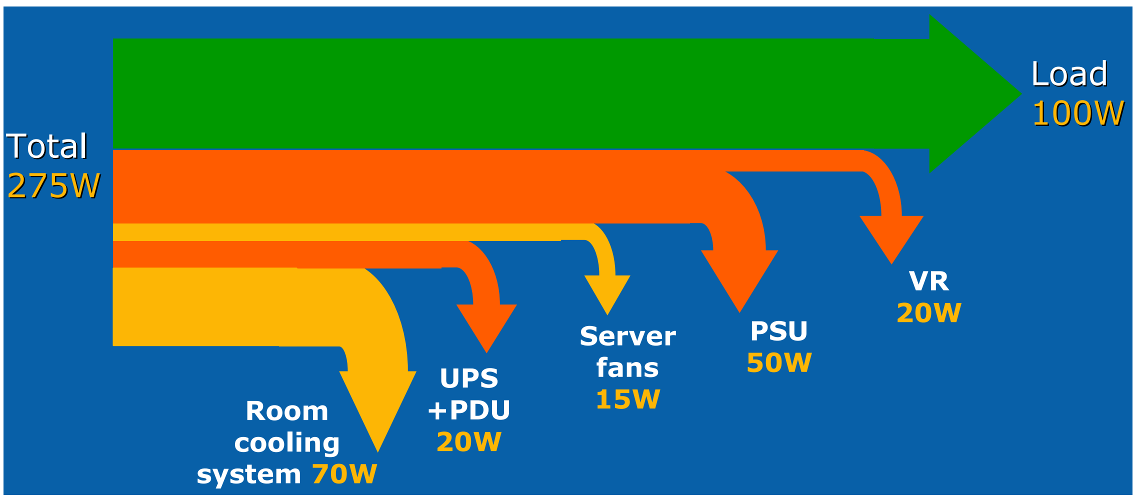

Source: Intel Corp.

The power losses due to inefficient power conversion from both outside and within IT equipment result in a large loss in useful electrical power, as well as directly increasing the energy required for HVAC systems to remove the heat produced in conditioned spaces. Figure 8 above details the conversion losses and the effects on cooling for a given representative computational load.

The above result indicates that for every watt of power utilized to process data, another 0.9W is required to support power conversion. In addition, for air-conditioned facility cooling, bench marking by LBNL suggests that another 0.6 to 1 watt (or more) of power will be required for each watt utilized to cool the power conversion equipment. While

[16] Additionally, the lower efficiency presented by the off-line power supplies utilized by data center equipment produces an extra load on batteries that only goes into producing heat instead of power conversion.

[17] Note that power conversion efficiency is limited by device technologies used, but include switching and conduction losses of MOSFETs, ESR of capacitors and inductors, etc.

additional cooling power can seem insignificant at the individual microprocessor level, when overall data processing activities reach power usage levels on the order of 200kW, or more (for example, a large Internet 4 hub installation) this can be a considerable load.

These trends provide an impetus to the industry to support research on better cooling methods and more efficient equipment/components, and this is happening on many fronts. But to look at the compute loads and determine ways to increase their efficiency or otherwise decrease their heat density, will also help alleviate the problem. In fact, from a cost savings point of view, it is even better, since every watt saved in a server means one less watt of cooling required.[18]

2.5 Demonstration Project Objectives

There have been many responses from the different sectors of the datacenter market to cope with this trend in rising power demand.[19] The responses vary in their focus, ranging from increased processor efficiency with demand management and/or multiple processor cores at the processor level, to liquid cooling at the server or rack level, and various HVAC strategies at the facility level. These efforts also include the CEC PIER Program-supported research to inform the IT industry, including our prior research on the efficiencies of server power supplies and UPSs, and EPA's ENERGY STAR® Programs. However, the efforts to date tend to focus on minimizing losses (or improving the efficiency of) specific components within the AC distribution system rather than optimization of the whole data center power delivery system.

An optimal system might integrate the IT equipment with the facility in such a way as to minimize power conversions. For example, the individual power supplies in servers could be eliminated if the correct voltages of DC power could be supplied efficiently from a central system, or in the case of fuel cells, directly from the power source. One industry expert envisions the data center of the future similar to a computer in its case. Taking this idea a step further, the electrical system could be thought of as an integrated system from where it enters the data center to the ultimate end use. When viewed in this manner, optimized systems could be designed so as to optimize energy (distribution and conversion losses), reliability, power quality, and potentially provide additional benefits such as elimination of harmful harmonics.

The concept of a DC data center is not new – most telephone companies' central offices use 48 VDC, and there are a number of DC datacenter proponents.[20] There are also

---

[18] However, given a more efficient server, datacenter operators would be more likely to increase the number of servers for a fixed cooling load rather than reduce the overall electric bill, because the demand for computing power is also rising.

[19] According to the Uptime Institute, the trend in power density for servers, storage devices and communications equipment will continue to rise, reaching new highs every year. From "How to Plan, Justify, and Manage a Major Data Center Project."

[20] Gross, P., and K. L. Godrich, "Total Integrated DC Datacenters." Paper presented at Intelec 2005.

companies offering DC servers today.[21] However, the current solutions on the market do not offer an integrated approach in the same way that AC datacenters are being designed and built currently. Although the DC solution is not new, unique or exceptional, it is still very difficult for most people to envision such a drastic change from the current norm, and it will take some time for the market to adopt the changes. Therefore, a demonstration of this technology using existing equipment (to the extent possible) will go far towards elevating the advantages of such an approach. In addition, demonstration projects tend to provide a critical step to verify engineering expectations for performance, and build awareness and confidence in the new technologies and practices.

The objectives of this Demonstration Project are to develop a power delivery system that does not contain as many power conversion stages using existing equipment and vendors where possible. One way to accomplish this system is to use DC power distribution. This project implemented a power delivery system that distributes DC to the server racks. Server manufacturers can supply products with power supplies that are ready to operate on DC.[22] The system used a single rectification stage, thereby removing the conventional UPS, transformer, and the rectifier in the server's first stage power supply. The removal of these stages, and the energy losses associated with these them, will achieve some energy savings automatically.

With a standard AC system installed next to this DC system, server loads were connected and programmed to run identical routines. For this identical amount of computing work, the input power for whole system was measured and compared. Taking all of the considerations, and considering the available resources, the project team identified the following goals for the demonstration project:

1. Show that DC-powered server(s) and/or server rack exists in the same form factor or can be built and operated from existing components with minimal effort.
2. Show that DC-powered server(s) and/or server rack provides the same level of functionality and computing performance when compared to similarly configured and operating server(s) (and/or server rack) containing AC power supplies.
3. Measure and document any efficiency gains from the elimination of multiple conversion steps in the delivery of DC power to the server hardware.
4. Identify areas requiring further development or follow up investigations.

[21] Humphreys, J. and J. Yang, "Server Innovations: Examining DC Power as an Alternative for Increasing Data Center Efficiency and Reliability." IDC White Paper sponsored by Rackable Systems Inc. August 2004.

[22] Servers used in the demonstration project were provided as "beta" equipment, which can be commercially available within six months with minimal certification effort

3. Project Chronology and Definition

As part of the server power supplies and UPS work, we identified a number of stakeholders in the latter part of 2005 and began dialogues with these industry representatives regarding a possible demonstration of an integrated DC power delivery system.

3.1 Timeline

Stakeholders Meeting – Fall 2005

The concept gained momentum through additional discussions and gained the CEC's support in the summer of 2005, an initial meeting for stakeholders was held in Chatsworth, CA in August of 2005, which was attended by over 20 industry representatives. The meeting was hosted by Pentadyne Power Corporation and Baldwin Technology, Inc.[23]

The meeting involved experts in data center design and operation, UPS vendor, controls manufacturers, technology integrators, instrumentation manufacturers, and component suppliers. No data center end users were present, but several were represented by the meeting participants. Several hours of lively discussion helped to identify potential industry partners, and brought out many issues and opportunities to address in the demonstration and eventual implementation.

During the discussion which followed, the results from prior investigations into UPS and power supply efficiency were presented and a number of considerations to be accounted for in a demonstration project were discussed. Barriers to implementing this technology were also introduced. Issues discussed include:

- All major server manufacturers currently have products that operate on 48V DC.[24]

- Commercially available equipment can be used for DC demonstration (no R&D required) including UL listed buss bar to supply the server equipment.

- The server manufacturers will always want to have control over the final conversion part of their equipment because of power quality/surge, safety, etc.

---

[23] The meeting also featured a demonstration of DC power distribution at the Pentadyne manufacturing facility. In this demonstration DC power was supplied at approximately 500 volts through commercially available components. The flywheel system maintains constant output at preset levels. The 500 VDC power is then converted to 48 VDC which can then be fed to the IT equipment. Depending upon the connected load, the flywheel will maintain this level for varying lengths of time. Baldwin Technologies supplied the DC delivery system, and Dranetz BMI supplied the monitoring instrumentation.

[24] Sun indicated at the meeting that it has swappable AC or DC power supplies.

- 6-7 kW per rack is the approximate limit for air cooling within a data center environment (although ASHRAE has shown an 18 kW rack in air (average today is 2-1/2 to 3 kW per rack).

- Data center user group sponsored by Emerson Network Power reports average of 60 W/square foot (this matches well with LBNL findings).[25]

- What should be the distribution voltage? 500V DC; 380V DC; 48V DC? Existing Sun power supplies can be changed out to run on 380V DC. IBM uses 350V DC. 380V DC was the consensus choice.

- Server manufacturers want their power source to be an integral part of the box – for safety, as well as to prevent others from plugging in.

- Barriers to DC power delivery included:
  - different voltages today in servers and other IT equipment
  - safety concern above 150V DC due to unfamiliarity
  - Above 600 volts - SELV classification
  - a new standard for voltage distribution may be needed

- Possible host sites were discussed

Kick-off meeting – April 2006

Discussions continued through the Fall of 2005 and into 2006, with additional manufacturers and industry representatives expressing interest and/or coming on board with commitment for participation. One outstanding issue crucial to the project implementation was finding suitable location(s) for the demonstration in California.

By Spring 2006, a suitable location for the demonstration project was found in Newark, CA, at Sun Microsystems' manufacturing facility in the San Francisco Area. With the location identified, this allowed the participant group to move forward with discussing the details for the demonstration. A meeting was held at this facility on April 3 and 4, 2006 for all interested parties in the demonstration. This meeting was intended to finalize the equipment list and scheduling for the tentatively named "DC Powering Architecture for Data Centers Demonstration."

A number of agreements were reached at this meeting, including:

> **Demonstration Configurations**: There was general agreement that both a rack-level and a facility-level demonstration could be accomplished, and needed for completeness and credibility.[26] Both could take place at the current Sun location

---

[25] The group thought this includes HVAC requirements.

[26] "Rack-level" is the term used to describe DC conversion at the server rack. This approach converts the facility's supplied AC into high-voltage DC via a rack-mounted rectifier unit. This approach concentrates the DC conversion into one unit/location in the rack, removing the AC to DC conversion function from the

if needed equipment could be found. For an “apples to apples” comparison, three different configurations were needed (one is needed as reference):

1. Configuration 1: Current data center typical set-up delivering 208/120V AC input to AC-powered servers (Reference Configuration).
2. Configuration 2: DC conversion/distribution at the rack level (Rack-Level Demo), using a rectifier unit to convert 208/120V AC at the rack, and delivering high-voltage DC to DC-powered servers.
3. Configuration 3: DC conversion/distribution at the building/data center level (Facility-Level Demo), converting 480V AC to 380V DC and delivering this directly to the DC-powered server units in the rack.

Equipment Set: Configurations 1 and 2 discussed above required additional equipment (UPS back up, PDU for power delivery) to accurately reflect the conditions at a typical data center, all three may also require additional loads to simulate operating conditions. The AC servers used would be matched to equivalent DC models and the same applications would run on both sets. Configurations 1 and 2 would share an AC buss way, while Configuration 3 would have its own DC buss way.

Testing and Measurements: The group agreed on test points and metrics, with emphasis on measuring the efficiency of the configurations – there would be no direct comparison of server equipment performance.

DC Input Voltage: Due to compatibility with existing equipment and devices, the group settled on 380 V DC for the high voltage DC input.[27] The test DC voltage would be set at 380 ± 5%, with 3% peak to peak for the ripple, which included high frequency and line frequency ripple components. However, the group agreed that this demonstration is about measurements rather than standard-setting. Other related points on DC input voltage were brought up by participants:

- Input DC voltage should be maintained between 300 volts and 400 volts for use with batteries and +/- 5% for use with flywheels.
- When using batteries, additional DC to DC conversion may be required.

DC Power Supply Efficiencies: Manufacturers noted that gains in the efficiencies of DC power supplies will probably be in the order of 1% to 3%. The server manufacturers already offer efficient DC power supplies at around 90%,

---

servers themselves. This approach can be attractive to operators who want to use DC powered servers, but not able to convert their facility exclusively to DC. “Facility-level” is the term used to describe DC conversion at the facility or data center level.

[27] HP and IBM were not available to present alternatives at this meeting (Intel noted that IBM currently uses 380 VDC in their high-end servers).

but end users can buy the less efficient (less than 80%) and less expensive power supplies (~10% less in cost). Thus for DC servers, the increase in on-board power supply efficiencies will come from the elimination of the computer on board power factor correction circuits and rectifiers. (Most of the efficiency gains were expected to come from the elimination of the AC to DC stages of the rack- or facility-level conversion).

Safety: The issue of high-voltage DC safety was brought up. The group agreed that, this was adequately addressed by the use of UL-rated equipment for high voltage DC, and other concerns were evaluated through a safety committee. For the future, we will research DC safety standards in use for other applications, including standards in use for the EU, military, and transit industries.

Other Related Technical Issues: The group also brought up various other issues, some of which were deferred:

- Construction Costs
- Reliability
- Equipment costs
- Grounding
- Environment
- Final Project Report

Equipment Demonstration Sessions: Sun Microsystems suggested that an "Open House" be held on 24 May 2006 to demonstrate the assembled equipment, with regular occurrences after that.

Code of Conduct: As the project was focused only on the demonstration aspect, participants also agreed to a "code of conduct" for cooperation and data sharing to minimize conflicts among competitors. A copy of this agreement is included in Appendix C.

Equipment assembly – May 2006

All participants at the April 06 meeting agreed that the months of April and May should be used for equipment configuration, definition, acquisition testing, and shipping, with the target date to finalize all equipment and transportation to the demonstration location was selected as 1 May 2006. On-site equipment assembly and verification took place in the first 2 weeks of May (1 May to 12 May), and testing/measurements took place on the 3rd week (15 May to 19 May).

Figure 9. Assembly at the demonstration site

In actuality, the gathering of equipment and shipping it to the demonstration facility took quite a bit longer, due to the complexity of the equipment, shipping times, as well as the need for approval of the proposed additional efforts. While a number of large pieces arrived in April and early May, it was not until the last weeks of the month that the assembly process actually started. Equipment assembly, integration, and testing took place in the last two weeks of May, and participants gathered for the first power-up on the first week of June.

Initial "Open House" – June 7, 2006

The first "Open House" was planned for May 24, and open to the press. However, a number of delays caused the date to slip, and with equipment "shakedown" still under way in the first week of June, the initial Open House was a more low-key event, with the press event rescheduled for two weeks later (June 21). Invitations were made only to interested industry representatives and colleagues.

Keeping the event low-key allowed the coordinating team of LBNL, Ecos, and EPRI Solutions to test presentation materials, go over equipment procedures, and trouble shoot any remaining issues prior to opening up the demonstration to a wider audience. The initial Open House also allowed the coordinating team to finalize the logistics and format of these events, which were used for the remainder of the summer.

Press Event: June 21, 2006

In order to achieve maximum media impact for the demonstration, one specific Open House was designated as a Press Event, to ensure that the project receives widespread industry attention. Participating industry representatives worked with their organization's PR staff to identify possible press contacts for invitation to the event. Both trade and popular publications, as well as electronic, television and radio were targeted. LBNL and the CEC and various industry partners' public relations representatives also participated in the process. LBNL's PR office prepared both the press invitation to the event and the press release for the day of the event, with input from other participants' PR representatives.[28]

Open House Events: July 12, 26; Aug 9, 16, 18

Participants agreed upon an every-other week Open House schedule, occurring on Wednesdays for the remaining of the summer, with adjustments for the July 4 holiday

[28] Electronic copies of the invitation and press release can be found at: http://hightech.lbl.gov/dc-powering/media.html.

and other competing events. The end date of the Open House was to coincide with the ending of the Phase 1 work for the LBNL team. Each of the Open House's two sessions (10 AM and 2 PM) averaged between 10 and 20 attendees from the data center industry and other sectors.

Due to unforeseen circumstances, the last Open House was changed to August 16, instead of August 23 as originally planned.[29] A last minute request from the US Armed Forces and a number of contractors resulted in a final, abbreviated Open House on the morning of August 18. After this the equipment set up was completely dismantled and returned to manufacturers or suppliers.

3.2 Participants

The demonstration project started with about 20 participants representing experts in data center design and operation, UPS vendors, controls manufacturers, technology integrators, instrumentation manufacturers, and component suppliers. As the project progressed and became known through a number of venues, the list of participants continued to grow throughout the duration of the project, new participants included data center end users, and a number of cellular carriers. In fact, industry representatives attending the "Open Houses" heeded our call for participation and a number quickly responded with additional equipment and/or in kind services from their respective industry sectors. Appendix A contains a full listing of the participants.

Participants of the Demonstration Project generally belonged in two groups:

- Contributors: Project "contributor" either contributed equipment, in-kind services, or both towards the project. A list of contributors is included below.

[29] The Newark demonstration facility was sold over the summer to a biotech firm, which was assumed control of the facility in September.

Alindeska Electrical Contractors, LLC
American Power Conversion, Inc. (APC)
Baldwin Technologies, Inc.
Cisco Systems
Cupertino Electric, Inc.
Dranetz-BMI
Emerson Network Power
Industrial Electric Manufacturing (IEM)
Intel Corporation
Nextek Power Systems
Panduit Corp.
Pentadyne Power Corporation
Rosendin Electric Inc.
SatCon Power Systems
Solara
Square D/Schneider Electric
Sun Microsystems
TDI Power
UNIVERSAL Electric Corp

- Stakeholders: Project "stakeholders" included industry experts and interested parties who regularly participated in discussions, provided feedback on approaches and other issues. Stakeholders did not contribute equipment to the project, but provided valuable insight into the process. A list of stakeholders and their industry sector is included below.

380voltsdc.com
CCG Facility Integration
Cingular Wireless
Dupont Fabros
EDG2, Inc.
EYP Mission Critical
Gannett
Hewlett Packard
Morrison Hershfield Corporation
NTT Facilities, Inc.
RTKL
SBC Global
Transistor Devices, Inc.
Verizon

3.3 Power Delivery: Definition

As discussed, the project participants helped to define the three configurations for the demonstration project. They are:

- AC Reference Configuration: This configuration is needed to simulate current data center typical set-up, delivering 208/120V AC input to AC-powered servers, and to be used as a reference to compare conversion efficiency.

- Facility-Level DC Configuration: This configuration is needed as the proof of concept – the ability to deliver high-voltage DC throughout the facility. This configuration handles the DC conversion/distribution at the building/data center level, converting 480V AC to 380V DC and delivering this directly to the DC-powered server units in the rack.

- Rack-Level DC Configuration: This configuration is needed to provide a possible migration option for AC data centers operators wishing to use DC equipment without facility-wide DC power distribution. This configuration accomplishes DC conversion/distribution at the rack level, using a rectifier unit to convert 208/120V AC at the rack, and delivering 380V DC to DC-powered servers.

3.4 Equipment Definition: AC Reference System

As detailed above, the AC reference configuration was intended to represent the typical set-up found in today's data center, taking 480 VAC from the utility feed, conditioning and stepping it down to 208/120V AC for input to AC-powered servers. In addition, due to uncertainty regarding manufacturers' abilities to provide sufficient servers to constitute a significant electrical load on the power delivery system, an AC load bank was added.[30] As configured, the set up has the ability to deliver from 15 kW up to 75 kW of electrical power to the server racks, while the projected actual load available for the demonstration was closer to 2 – 4 kW, or about 4% to 8% of the available system power. The load bank would allow the simulation of 30 kW to 40 kW power demand on the system, more accurately reflecting the loading of current and future server racks in data centers. The simplified AC power delivery system is shown below.

Figure 10. AC Reference Configuration

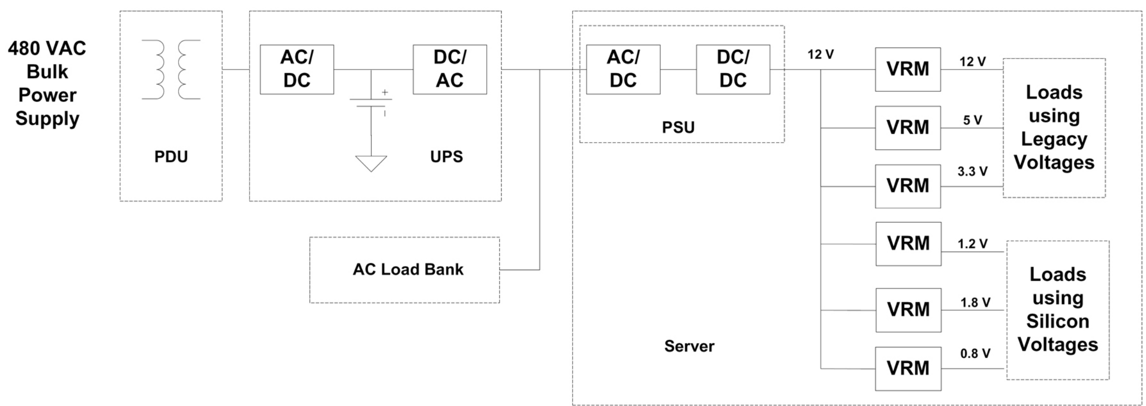

In this diagram the UPS is a double-conversion UPS (typical), preceded by a transformer in a power distribution unit (PDU). Inside the server is an AC/DC conversion at the input to the server power supply. There are still many DC/DC converters at lower voltage levels, known as voltage regulator modules, or VRMs, but those are not the focus of this project. So to implement this system, the complete equipment is listed below:

- Double conversion UPS
- Transformer (for PDU)
- Servers
- Supplemental load banks (to load the UPS to a typical level)
- 600 V rated busway

[30] This was a commercially available product, leased by the Demonstration Project expressly for this purpose.

- Various circuit breakers
- Wire and other hardware

3.5 Equipment Definition: Facility-Level

The facility-level DC configuration was intended to replace the typical AC set-up found in today's data center, taking 480 VAC from the utility feed, convert it to DC voltage, and deliver the DC voltage to the racks for input to DC-powered servers. This configuration is needed as the proof of concept – demonstrating the ability to deliver high-voltage DC throughout a data center facility. As in the AC set up, the uncertainty regarding manufacturers' abilities to provide sufficient DC servers to constitute a significant electrical load on the power delivery system, a DC load bank was added.[31] The DC load bank would also allow the simulation of 15 kW to 40 kW power demand on the system, more accurately reflecting the loading of current and future server racks in data centers.

The next diagram shows the arrangement for DC power delivery to the data center loads.

Figure 11. Facility-Level DC Configuration

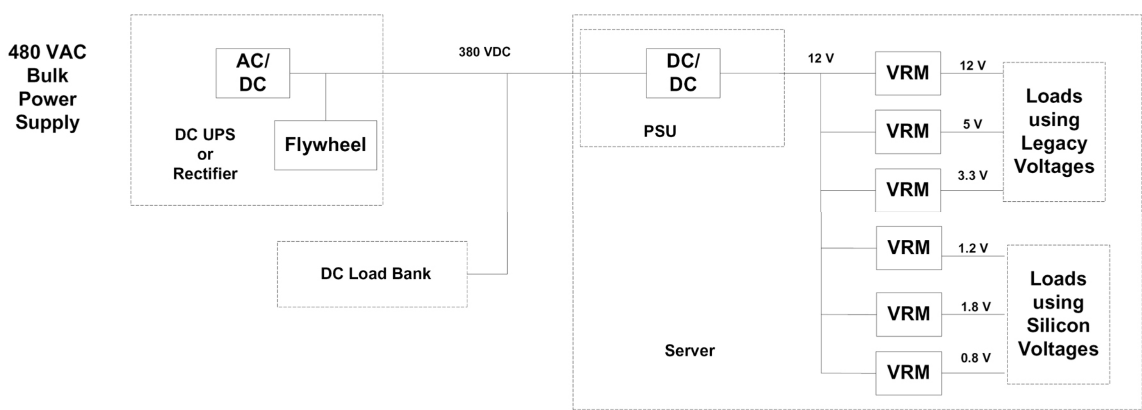

In this layout the only power conversion before the server takes place through one rectifier. Inside the server power supply, only one DC/DC conversion occurs. In addition to the reduction in losses, this arrangement is inherently simpler, which will should lead to lower equipment cost, lower installation cost, and higher system reliability due to lower parts count. The necessary equipment for implementation is listed below:

- System level rectifier
- Servers, modified to take 380 VDC
- Supplemental load banks (to load the UPS to a typical level)

[31] This was a custom built DC load bank, designed and constructed specifically for the use of this Demonstration Project.

- 600 V rated busway
- Various DC circuit breakers
- Wire and other hardware

3.6 Equipment Definition: Rack-Level

The rack level demonstration was designed to show that 380 VDC could be supplied at the rack level and distributed throughout the rack. Instead of using a rectifier at each power supply, only one rectifier would be necessary to convert AC to DC for the entire rack. The advantage for this approach would be to change the heat removal requirements – it may be more economical to remove the heat at the rack level rather than from the servers themselves. Measurement of cooling effectiveness was not in the scope of this project; rather, this arrangement was intended to show feasibility of design. In addition, this configuration is needed to demonstrate that a possible migration option is available for AC data centers operators wishing to use DC equipment without facility-wide DC power distribution.

Figure 12. Rack-Level DC Configuration

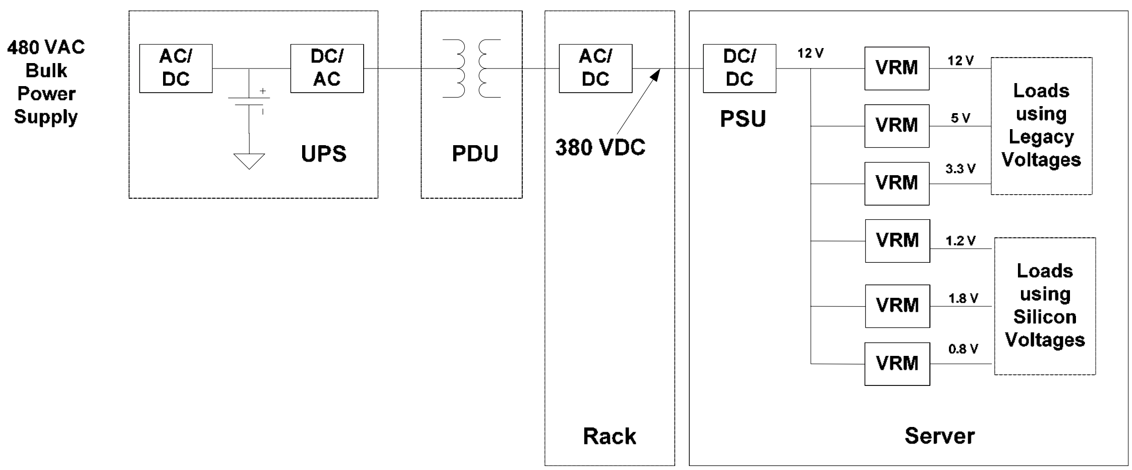

So to implement this system, the complete equipment is listed below:

- Double conversion UPS
- Transformer (for PDU)
- Rack level rectifier
- Servers, modified to take 380 VDC
- Supplemental load banks (to load the UPS to a typical level)
- 600 V rated busway
- Various circuit breakers

- Wire and other hardware

3.7 Equipment Definition: Monitoring

The monitoring system used has high bandwidth data capture capability using voltage probes and current transformers (CTs). The system uses a central node called an info node and multiple individual modules called data nodes that reside close to the point of measurement. The system has the capability to capture waveforms during power line disturbances, and also to capture steady state data. In this case, the system was used to capture steady state voltage, current and power data. Figure 13 below illustrates the layout used by this system. In addition to providing real-time data measurement to the system operators, this system has the ability to provide "live" data display via the internet. In fact, measurements were made available via the LBNL website for the duration of the demonstration project.

Figure 13. Diagram of the Monitoring System

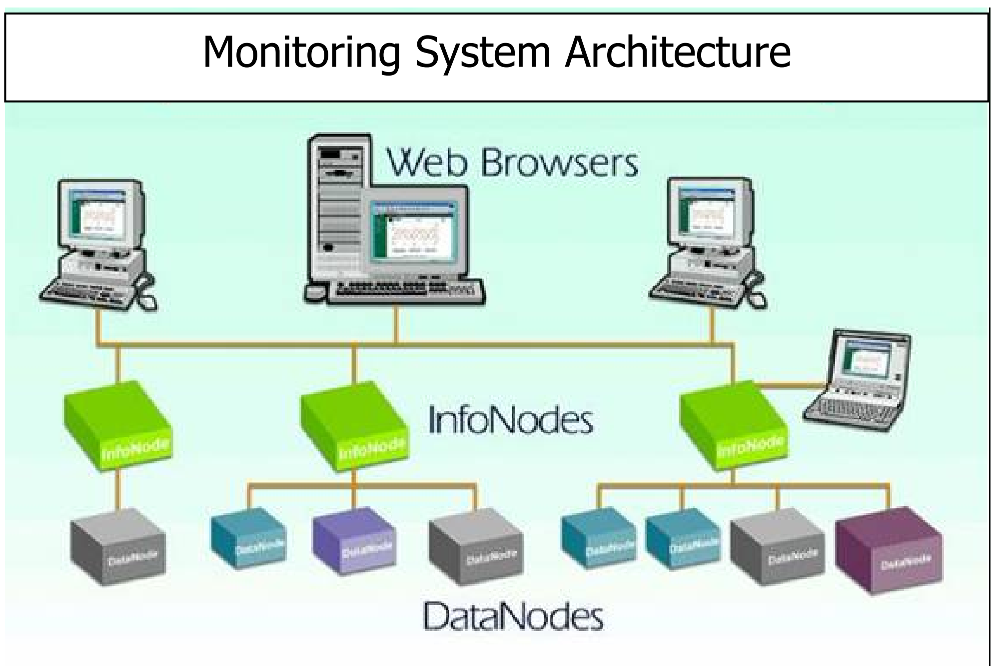

The placement of the monitors can be seen in the one-line diagrams below (a full-page layout is referenced in Appendix B).

Figure 14. AC Distribution Layout and Monitoring Points

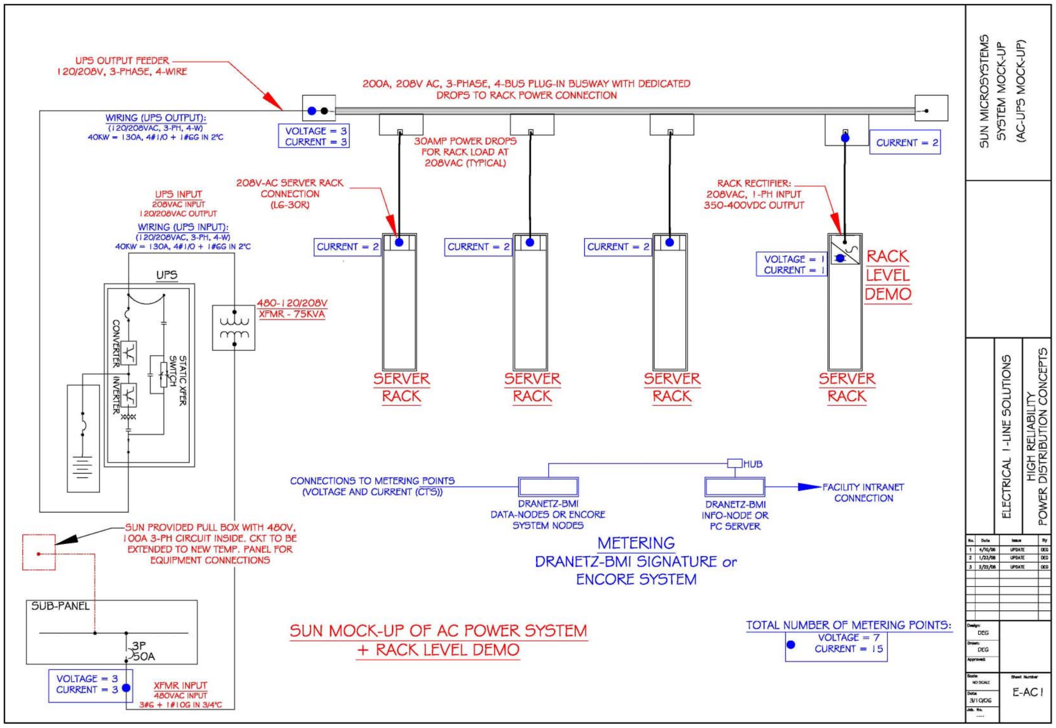

Source: Baldwin Technologies, Inc.

The above diagram shows the AC Reference System and its monitoring points, along with the USP system's batteries and transformer. The system distributes 208 VAC to the server racks through 30 A breakers. Also shown is the rack level demo. The Monitoring system is shown as blue dots with callouts.

There are 7 voltages and 15 currents measured.

Figure 15. DC Distribution Layout and Monitoring Points

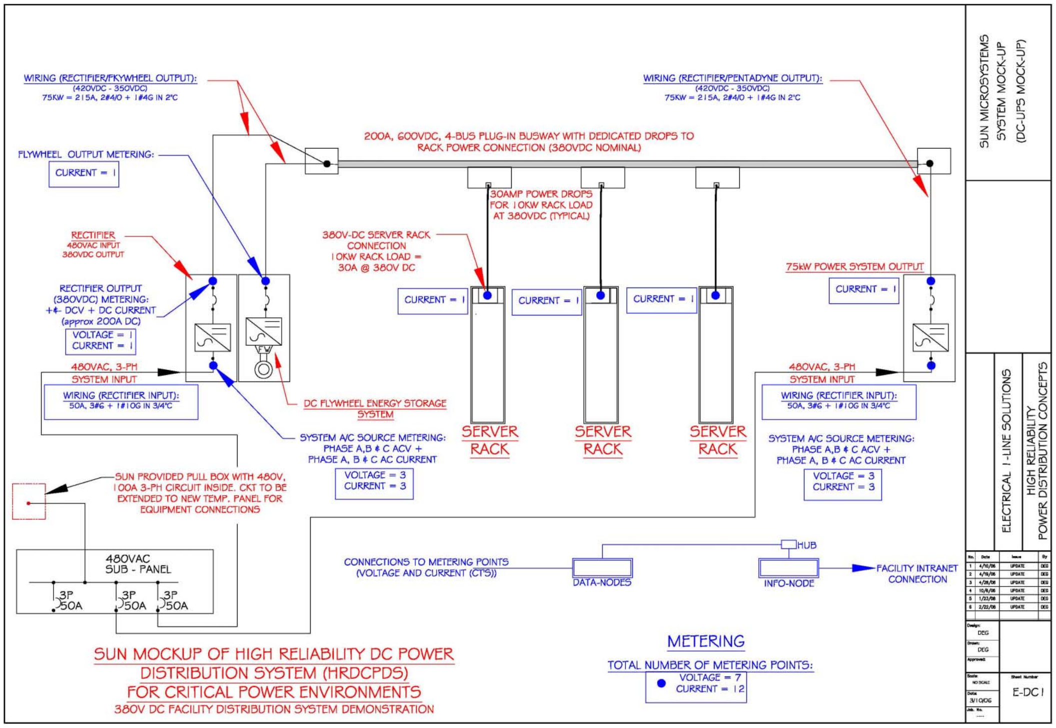

Source: Baldwin Technologies, Inc.

The diagram above shows the monitoring points for the DC Demonstration set up. Note that two main rectifiers are shown in this diagram. Only one is necessary, but a second one provides redundancy and a second vendor was able to participate. A flywheel is used for energy storage (ride through for interruptions), but batteries can also be used with DC distribution. The DC bus is energized at 380 VDC, with individual racks powered through 30 A breakers on the buss way.

The monitoring points are shown on this diagram as blue dots with callouts. There are 7 voltages and 12 currents shown.

3.8 Other DC Systems Demonstrated (48V DC)

The interest generated by the project helped to bring about other manufacturers of DC equipment, notably manufacturers of communication and support equipment using 48 V DC. These manufacturers provided additional computing, telecommunication equipment and racks for the use of the demonstration project. The presence of these manufacturers and their equipment provided another effective reminder that DC power distribution is not a new concept, and has been safely and effectively used in telecommunication and data networks. In addition, it also demonstrated that DC and AC power delivery system can co-locate in a data center facility.

The project coordinating team had initially also investigated the issue of 48V DC power delivery systems and related efficiencies, using available data from manufacturers and other industry experts.[32] This supplemental work was intended to inform the team on selecting an appropriate DC voltage for the demonstration. However, input from the participants, especially the server manufacturers, helped narrow down the demonstration voltage to 380 V DC. This obviated the need for an in-depth comparison of different DC voltages.

These supplemental DC systems were included in the overall demonstration for completeness only. Due to limited resources, monitoring equipment constraints, timeline, and priority for implementing the rack- and facility-level demonstration configurations as well as obtaining measurement, these 48 V DC systems were not monitored or their efficiencies measured. Without comparable AC equipment, it was not possible to carry out the same types of comparison that was being performed on the other configurations.

[32] Gross, P and Godrich, K. L. "Total DC Integrated Data Centers." Paper presented at 2005 Intelec.

4. Implementation Issues

As with any prototyping projects, where discrete equipment and products from different manufacturers have to work together for the first time, there were a number of challenges. In discussion regarding implementation within the group, we had agreed to divide the configurations into two distinct areas: outside of the server rack, and inside the server rack. This allowed the server manufacturers to focus on the DC conversion and adaptation for their products to be used in the server racks, while the facility level team focused on integrating the power delivery infrastructure.

Actual implementation and integration of the demonstration configurations were accomplished in an astonishingly short time, thanks in part to the commitment and enthusiasm of the participants. In fact, a number of industry "veterans" remarked that this Demonstration Project was one of the quickest ones that they had experienced. Although everything came together well, quickly, and every effort was made to use available, off the shelf components, there were a number of issues that we encountered that were not entirely solvable with existing equipment, or needed further investigation. They are outlined below.

4.1 Commercially available equipment/missing components

One of the stated objectives of the project was to accomplish the demonstration with commercially available products to the extent possible, especially with the power delivery infrastructure. The project was mostly successful in this regard with the following exceptions:

Within the server racks, we identified a number of components that are not yet commercially available:

1. Servers: There currently are no servers commercially available for 380V DC operation.[33]

2. Power cord/Power plug or connector: Although there are commercially available power cord and plug/connectors for 380V DC. To enable widespread adoption of this technology, the industry will need to standardize on only one type of connector (similar to an AC cord).[34]

3. Power "strip": There are there no 380 V DC powers "strips" for power distribution within a server rack. Similar to the power connector, it is desirable for the industry to standardize.

---

[33] Equipment used is not yet safety and EMI certified.

[34] We were able to locate a DC connector after the implementation had taken place, but it is not a universally accepted product. Note that the demonstration project used standard AC connectors and power plugs, but this may have serious implications in a mixed voltage environment, such as hot swapping servers.

Outside of the server racks, we also encountered a number of issues related to lack of commercially available equipment, including

Figure 16. DC Power Drops

1. Power drops/connectors: There are not yet specifically designated connectors for powering server racks. (The team used standard Hubbell connectors and standard APC equipment racks).

2. Circuit breakers: The availability on the market of DC circuit breakers specifically for this application was scarce. In fact, UL listed breakers were one of the more difficult items to obtain for integration.

4.2 Grounding, Bonding and Protection

Grounding, bonding and protection are issues with important implications for the adoption of DC power systems. These received significant discussion at the Open House sessions, and warranted additional research by the team. We identified a number of documents suitable for DC application (although not specifically for 380V DC), which are summarized in Table 3 below.

These industry documents are for the most part accounted for and harmonized to the fullest extent practicable in the IEEE Emerald Book. Formally titled: *IEEE Standard. 1100-2005 – Recommended practice for Powering and Grounding Electronic Equipment,* the Emerald Book provides a useful general reference in terms of the grounding and bonding recommended practices for facilities that simultaneously contain both AC power systems and DC power systems.

The Emerald Book integrates many of the traditional telecommunications recommendations and discusses how to integrate the AC and DC power systems to accomplish the important safety and performance objectives of each. The key concepts related to grounding and bonding have to do with the connection to the buildings grounding electrode system which the Emerald Book refers to as the "common bonding network and when to use a single point of connection to the grounding electrode system to form an "isolated bonding network." These objectives and the concepts employed are detailed in the following sections along with some historical background.

Much of the guidance on grounding and bonding DC power systems for telecommunications and information technology equipment is rooted in the traditional telephone (telecommunications) utility (regulated) industry. The basis of this guidance is supported in IEEE Std. 1100-2005 for the commercial (deregulated) industry - with some modifications made to meet requirements of the commercial market segment.

Table 3: Industry Documents Associated with DC Applications

ANSI	T1.311-1998	DC Power Systems—Telecommunications Environment Protection
ANSI	T1.315-1994	Voltage Levels for DC-Powered Equipment Used in the Telecommunications Environment
ANSI	T1.330-1997 (R2002)	Valve-Regulated Lead-Acid Batteries Used in the Telecommunications Environment
ANSI	T1.333-2001	Grounding and Bonding of Telecommunications Equipment
ANSI	T1/TIA J-STD-607-A-2002	Commercial Building Grounding (Earthing) and Bonding Requirements for Telecommunications
IEEE	Std. 1100-2005	IEEE Recommended Practice for Powering and Grounding Electronic Equipment
NFPA	70-2005	National Electrical Code (NEC®)
NFPA	75-2003	Standard for the Protection of Information Technology Equipment
Telcordia	GR-1089-CORE 2006	Electromagnetic Compatibility and Electrical Safety - Generic Criteria for Network Telecommunications Equipment
Telcordia	GR-295-CORE 2004	Isolated and Mesh Bonding Networks: Definition and Application to Telephone Central Offices
T1A	TIA/T1 J-STD-607-A-2002	Commercial Building Grounding (Earthing) and Bonding Requirements for Telecommunications
T1A	TIA 942	Telecommunications Infrastructure Standard for Data Centers
UL	UL 60950-1 - Edition 1 04/03	Information Technology Equipment - Safety - Part 1: General Requirements

Source: Panduit Corp.

Historically, telecom is DC powered and information technology equipment has been AC powered. Therefore, the grounding and bonding standards for these different types of power systems evolved differently due to the DC being predominantly utilized in a regulated environment considered “under the exclusive control of the utility” (NFPA 70). As such, DC power system operating voltages are affected by factors such as:

1. Battery technology (vented or valve-regulated)
2. Rectifier output regulation in maintaining constant voltage under dynamic loads
3. Voltage drops in DC power conductors
4. Operating voltage limits of various connected loads
5. Derating for environmental conditions such as altitude

4.2.1 Telecom DC Systems Grounding Example (up to 160V DC)

Generally, telecom DC power systems date back to the first telephone systems where DC was used for talk battery, signaling circuits and for operating switching and control relays. Large centralized (bulk) DC power plants (systems) had primary components such as rectifiers, power-boards, primary and secondary distribution feeders, and fuse bays. These were built with reliability in mind, and space and efficiency were secondary considerations. The DC power system was grounded to earth (often more than once). The grounded conductor was termed the “Return.”

Modern DC power systems are more compact, use much smaller footprint components, and are more efficient. Today, a small centralized DC power system can be contained in a single rack. For standards compliance purposes, this equipment is generally classified as information technology equipment (ITE). For the purposes of evaluating DC powered ITE, robustness of the system is a key consideration. Accordingly, the design of grounding, bonding and protection of the installation is based on specific system considerations.[35]

Generally, the modern centralized DC power system is designed to operate as a single-point grounded system. An equipment-grounding (bonding) conductor is installed as part of the distribution circuit to ground any metal parts of the ITE and to clear any ground faults by facilitating the timely operation of the upstream overcurrent protection device. The system grounding (electrode) conductor is termed the DCG and is connected to the Return near the DC power source. The Return is only grounded once.

This arrangement is extremely similar to grounding and bonding an AC power system per NFPA 70-2005 – with some possible variations allowed by UL60950.

Figure 17: Similarity of Recommended Grounding for AC and DC Power Systems and Load Equipment

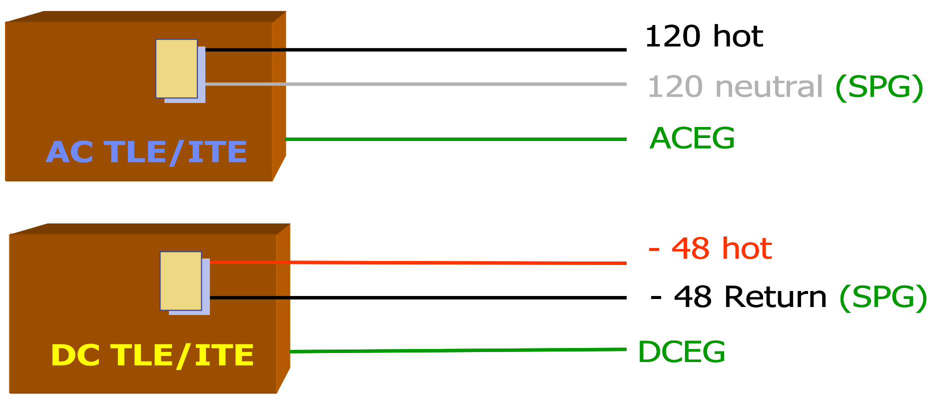

Source: Panduit Corp.

[35] For example, these considerations include:
1. Is the ITE suitably robust (per Telcordia GR-1089-2006) to operate in a Common Bonding Network (CBN)? Or must it be operated in an Isolated Bonding Network (IBN)?
Note: The significant role of the IBN topology is to isolate the ITE from currents expected flowing through the CBN – especially lightning. Availability requirements for ITE in a data center may not need to be specified to the "high-nines" required for a telecommunications service provider (TSP) – such as at a telecommunications central office. The availability specification will determine if an IBN may be appropriate for the data center – or a portion of the data center.
3.Is the DC power supply dedicated to a single type of equipment bonding network (CBN or IBN) or is it to be a shared resource?
4.Where is the planned location for the DC power supply relative to the location of the ITE?
Is any of the ITE required to integrate the Return and DC equipment grounding conductor (DCEG) at the DC power input?

Figure 18. DC Power System Showing Single-Point Grounded Return[36]

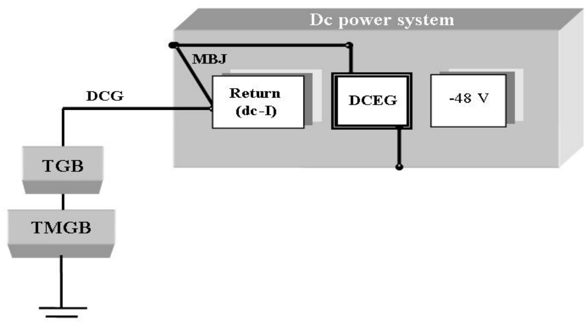

Source: William Bush

Figure 19. Grounding of DC Power Input for IT Equipment[37]

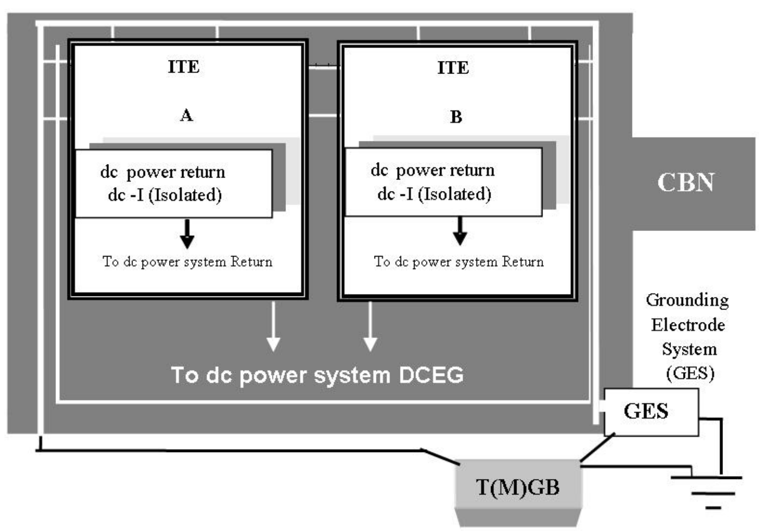

Source: William Bush

[36] SPG = Single point ground; ACEG = AC equipment grounding conductor; DCEG = DC equipment grounding conductor; TLE/ITE = Telecommunications Legacy Equipment / Information Technology Equipment

[37] DCEG = DC Equipment Grounding Conductor; TGB = Telecommunications Grounding Busbar; TMGB = Telecommunications Main Grounding Busbar; MBJ = Main Bonding Jumper; DCG = DC System Grounding Conductor.

Figure 20. Illustration of the Common Bonding Network (CBN)

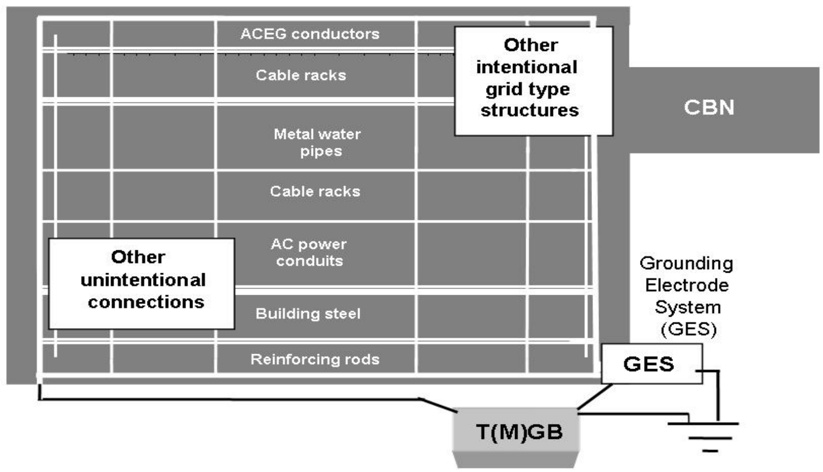

Source: William Bush

Figure 21. Illustration of the Isolated (insulated) Bonding Network

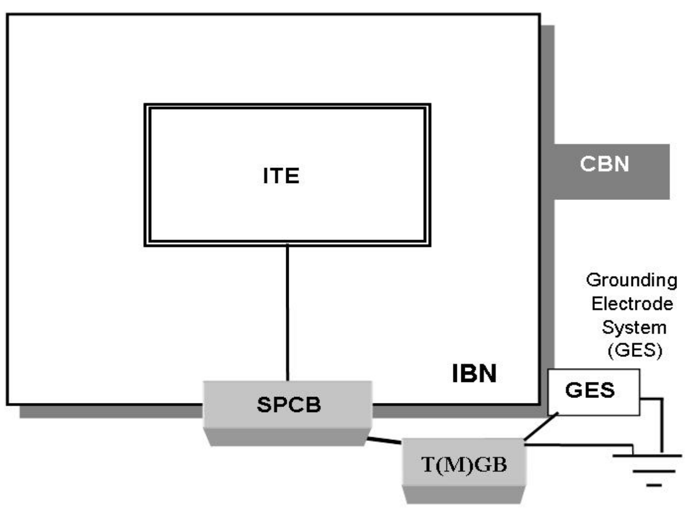

Source: William Bush

4.2.2 Additional considerations

Porting the DC power system telecom utility practices to the non-regulated environment requires additional considerations, they include:

1. The highest DC voltage covered by the telephone/telecom/ITE industry is 160 (ANSI T1.311). The utilized 380V DC (or higher) utilized by the demonstration project is essentially new territory and will require some additional safety and performance investigation. However, certain established principles for DC below are not expected to change.
    a. Over-current protection devices (OPD) and disconnect devices for DC power systems will need further investigation for the greater than 160V DC systems. AC rated devices do not automatically port over to the DC power system at the same rated voltage and current. The established 2:1 protection coordination scheme for DC fuses is not readily applicable since data centers typically utilize circuit breaker technology.
    b. Transients for other than 48V DC systems are not well described – if at all.
    c. Battery disconnects (where batteries are used as standby resource) considerations: 1 pole or 2 poles, and their interrupt ratings
    d. Conductor sizing is not a 100% straightforward utilization of the several Tables usually utilized from NFPA 70-2005 for AC conductors.
2. The rectifier technology is assumed to be switched mode power supply (SMPS) which may involve noise control via power line filters and additional audible noise from cooling fans.
3. Rectifier operation modes at the higher voltages such as 380V DC may need verification for ac input power factor correction, load sharing, paralleling, voltage sensing, current limitation, etc.
4. The DC power distribution topology is historically buss or cabling. The buss version described in this project will likely not be the only topology utilized. Further, there may also be a need for an under-floor topology. Distribution (such as rigid busbar) to withstand under extreme fault current conditions will likely to vary considerably from that for a telecom 48V DC system.
5. Design parameter for voltage drop from standby resource (such as a battery plant) to the load equipment. For telecom, the parameter is typically 1 V per 24 Volts of the supply. 48 V supply = 2 V drop.
6. Battery sizing: For a 380V DC battery plant, the size of the battery rack (if metallic) grounding conductor needs to be considered.
7. Centralized DC power systems are restrained to use in a Restricted Access Area (RAA) by UL 60950. Note: The data center is considered a RAA.
8. Grounding: Should the 380V DC system be operated as a positive grounded system (48V DC telecom) or negative grounded system (24V DC telecom).

9. The location of attachment of the DC system grounding electrode conductor (DCG) to the Centralized DC power system is allowed to occur forward from the source toward the telecom load equipment per UL60950.
10. The DC power system equipment grounding conductor (DCEG) is allowed to be routed separate from the DC supply and Return circuit conductors per UL 60950.
11. The DCEG is permitted to be bonded to the Return at the load equipment per UL 60950.
12. Considerations 9-11 reflect telephone company practices accommodated by first UL 1459 and then UL 60950.
13. The CDCPS historically can serve two recognized equipment bonding topologies.
    a. Common Bonding Network (CBN)
    b. Isolated (Insulated) Bonding Network (IBN)
14. The CBN and IBN equipment bonding topologies were described in IEEE Std. 1100-1999. TIA 942 also notes the CBN topology directly and the IBN through reference to IEEE Std. 1100-2005.
15. Based upon these considerations, the prudent approach is to utilize IEEE Std. 1100-2005 as the base document for grounding and bonding the DC power system in a data center.
    a. Essentially mirrors topology for an AC power system
    b. Single-point grounding of the DC power system at the source location
    c. A co-routed DC equipment grounding conductor with the circuit wiring (supply such as -48 V, Return)
    d. Bonding of DC equipment grounding conductor to the "Return" at the load equipment is prohibited
    e. Fully controlled paths for direct current

4.2.3 Supplemental grounding and bonding

For the data center, a telecommunications grounding and bonding infrastructure in accordance with J-STD-607-A, TIA 942 and IEEE Std. 1100-2005 is expected.[38] This infrastructure is bonded to the electrical power grounding electrode system, to building steel (where accessible) and to the serving AC power panel equipment ground at each floor. Grounding needed for the data center equipment(s) is obtained from the expected ground bar on that floor (such as a Telecommunications Ground Bar – TGB). Note that this grounding and bonding infrastructure is not the same physical structure as the grounding infrastructure that might be placed for the electrical power system.

[38] In no situation should a totally separate grounding system be deployed as this will bring to bear both safety and performance problems.

Figure 22. Illustration of Telecommunications Grounding and Bonding Infrastructure

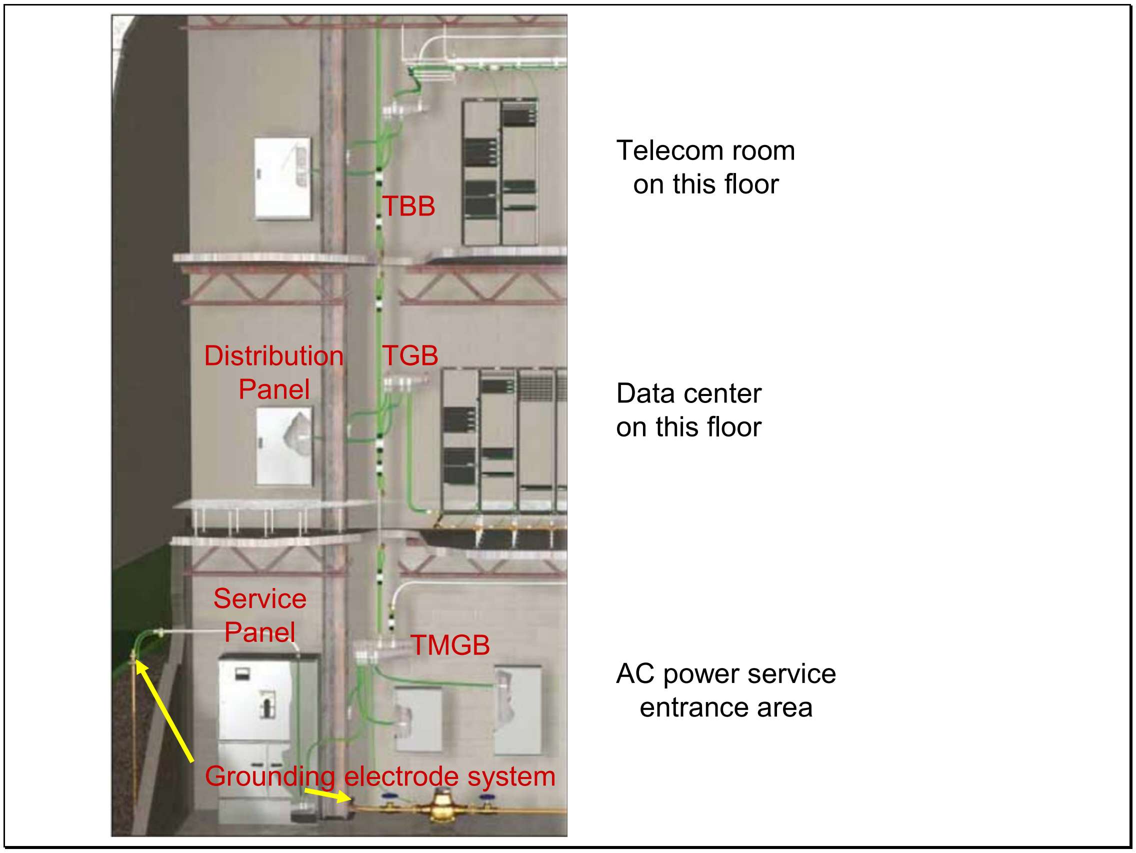

Source: Panduit Corp.

Noise concerns for the data center equipment do involve common mode noise generated and distributed by the power system to the electronic load equipment. Generally, the equipment AC input power supplies are quite tolerant of common mode noise. It should be expected that the server and other manufacturers will design and test equipment (DC input power supplies) to be similarly robust. [Already accomplished for 48V DC telecom equipment meeting the requirements of Telcordia GR-1089-CORE-2006 and placed into a Common Bonding Network (CBN).]

A supplemental signal reference structure (such as an under floor signal reference grid), where properly designed, installed and utilized) can be effective for noise control across low-frequency single-ended signaling links or poorly designed communication links. However, in many cases the design is limited to a few MHz (at best) and severely compromised by large loop areas between the cabling and the reference structure. Modern high bandwidth cabling (especially fiber optics) is significantly robust and the added value of a signal reference grid is very questionable. IEEE Std. 1100-2005 recommends "decoupling" of links for distributed systems since distributed "common" grounding systems are suspect after low-to-mid frequencies and often impractical after several meters.

As noted in IEEE Std. 1100-2005, the default equipment bonding topology is the Common Bonding Network (CBN). The CBN can be readily utilized for efficient direct bonding of equipment and other apparatus to the grounding system. Such an arrangement provides efficient grounding and inter/intra-unit bonding of metal cabinets, racks and miscellaneous metal objects (especially when they are not powered). Electrostatic charge buildup and dissipation is also greatly aided by the CBN multiple grounding paths.

Figure 23. Data Center Grounding Infrastructure at the Room Level

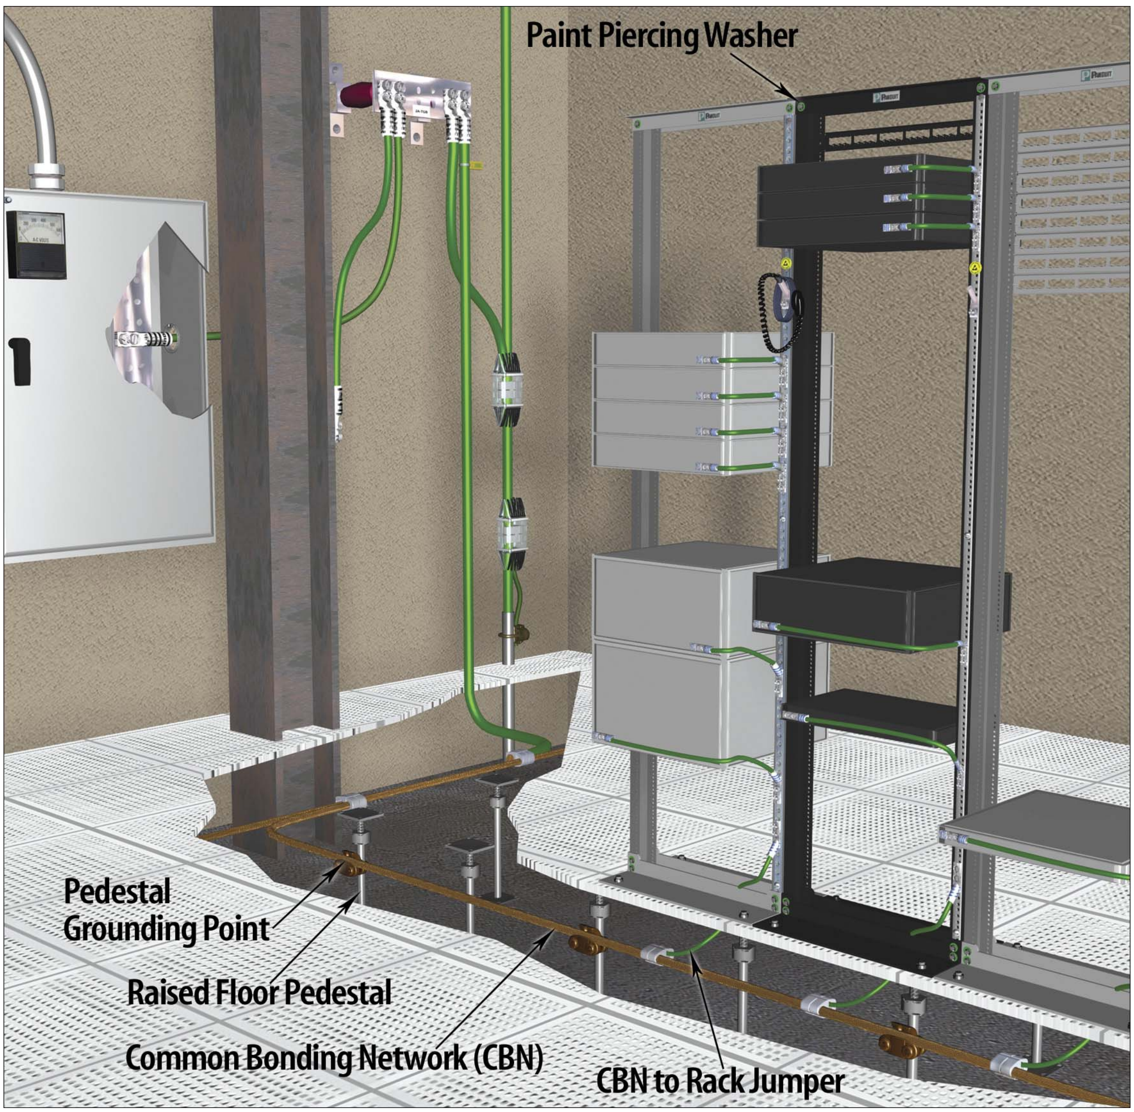

Source: Panduit Corp.

Based upon these considerations, the prudent approach is to utilize IEEE Std. 1100-2005 as the base document for supplemental grounding and bonding of the data center. This approach has the following advantages:

1. Harmonized with J-Std-607-A and ANSI T1 .333
2. Compatible with NFPA 70 and NFPA 75
3. Promotes the use of suitably immune links or decoupled links (fiber optic link is ultimate decoupling method)
4. Promotes the use of equipment with recognized immunity (less susceptibility)
5. Promotes the recognition and proper application of certain equipment bonding topologies (variations of CBN and IBN)
6. Default supplemental grounding is an extended CBN such as an above rack or under floor conductor grid.
    a. Not intended to function as a signal reference grid
7. Where otherwise deemed a requirement [historically unlikely unless for older generation Electronic data process (information technology) equipment], a signal reference grid suitable for up to a few MHz is also described.

4.3 Reliability

We received a number of inquiries from open house visitors about the DC delivery systems' reliability, and this topic certainly warranted additional investigation. There are a number of existing studies covering this topic; however, they all deal with the reliability of the telecom industry's 48V DC systems. We found no available studies dealing with the reliability of high-voltage DC (as compared to 48V DC) systems.

However, there are a number of similarities between the 48V DC systems and our Demonstration set up, and the following inferences can be drawn regarding DC power delivery in general:

- Ability to shift large heat loads out of the servers.
- Reduction of overall power consumption for each server box.
- Reduction in overall component count (from utility to silicon).
- Higher DC UPS reliability record.[39]
- Reduction of power conversion steps.

All of the available studies indicated that DC power distribution has a higher reliability record than AC distribution, often by several orders of magnitude, due to the above factors. NTT Facilities, the largest operator of data centers in Japan with close to ~80%

---

[39] Åkerlund , J. "DC Powering of Internet Certifies for Telephony." Uninterruptible Power Networks, UPN AB, Sweden, 2004.

of the market, also has similar data supporting increased DC power delivery reliability that they have collected in support of DC power distribution. While a definitive conclusion cannot yet be made regarding DC system reliability without additional data, it seems that a DC approach would not reduce data centers' reliability in any way.

4.4 Costs

Another issue that received a considerable amount of attention and generated considerable discussion is the cost of implementing DC power delivery systems. Unlike the available studies on DC reliability, there is less information on the costs of DC data centers. A detailed analysis of cost was outside the scope of this project, and we were not able to directly compare costs due to the fact that some of the components were not directly available, or are only available on the market in small quantities, which significantly affect costs.

Nevertheless, the project team was able to collect a number of anecdotal evidence on costs, presented below:

According to a recent white paper, the cost of a DC power plant is typically lower than an AC UPS system by 10% to 30%.[40] However, this source also claims that the additional engineering and distribution costs, along with the need to power some AC-only equipment can increase costs of a DC system for a data center.[41] In addition, manufacturers have confirmed that DC servers and storage array costs will likely be higher, at least in the near term due to the lower demand volume.

The sources that we consulted on costs all agree on one area where DC data center costs will likely be lower: TCO, or total costs of operation. This is due to a number of factors, including lower energy costs (from a reduction in both conversion - heat generation and HVAC – heat removal), lower replacement costs (from higher equipment reliability). As about half of total IT costs are related to hardware and software, with the rest in administration and overhead costs, the lower power and maintenance costs also help to reduce the administration costs.[42]

Our preliminary investigation shows that data is available for AC power delivery system costs, so that a related future investigations should consider a direct comparison of costs, from the design, build, and equip to TCO for both AC and DC data centers. Of particular interest would be a comparison of the costs to build and operate AC and DC data centers

[40] "AC vs. DC for Data Centers and Network Rooms" APC White Paper #63, American Power Conversion, 2003.

[41] Note that for our Demonstration project, the DC conversion and distribution was specifically designed to deliver DC voltage to the equipment racks as well as lighting.

[42] Burt, J. "Can DC Power Cut Data Center Costs?" *eWeek*, March 27, 2006. "Thermal Management & Server Density: Critical Issues for Today's Data Center," White Paper, Rackable Systems, 2004.

of a typical size, such as a 5 MW or 10 MW center. Coupled with reliability data, this information will go far to help inform decision makers on future data center designs.

4.5 Measurement Verification

For the Demonstration Project, two sets of measurement were made, one with the current transducers or transformers (CTs) supplied with the Signature System, and one with a commercially available Power Analyzer. This was done to ensure the reliability of the measurements, and to address the fact that one of the potential issues with the demonstration configuration could be the accuracy of the measuring instruments.[43] The approach used for the Demonstration Project is briefly outlined below.[44]

The measurement setups for AC and DC measurements respectively are shown in Figure 24. At each measurement point, voltages and currents are measured and converted to digital outputs. The setups for AC and DC measurements are shown in Figure 24. For both AC and DC, the voltage is applied directly to the voltage measurement pod, which outputs a proportional voltage signal to the data node.

Figure 24. Measurement Points

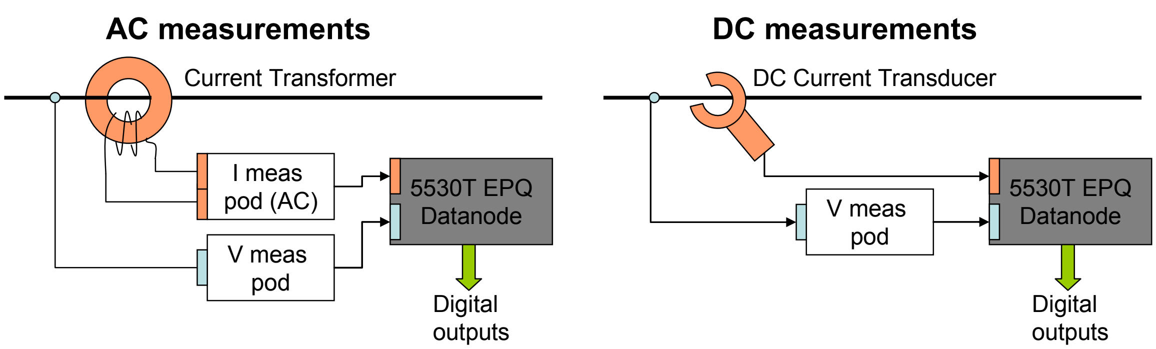

For AC measurements, the current transformer generates a current proportional to the primary current. This current is an input to the current measurement pod, which outputs a voltage to the data node where the analog voltage is converted to a digital output. For DC current measurement, the current transducer outputs a voltage proportional to the current, and this voltage is input directly to the data node. Table 4 lists the measurement accuracies for each piece of measurement equipment used.

---

43 DC current transducers with extremely high accuracy (< 1%) can be harder to locate, especially under the time constraint of the implementation schedule and the quantities required, so standard DC CTs, (~1% accuracy) were used.

44 Analysis prepared by A. Pratt (Intel), and R. Ignall (Dranetz BMI), August 4, 2006.

Table 4. Measurement Equipment Accuracies

Description	Part Number	Range	Accuracy	Comments
Data node and measurement pod	5530T EPQ Series		±0.1%	Accuracy spec includes both data node and pods
AC Current transformer	AL-151		±1%	150A full range
DC Current transducer	PR150/SP1	15 to 180 A DC	±1%	150A full range
Server input power			±0.3%	

An error is introduced at each step in the measurement process, and the overall measurement accuracy is a function of all these errors. The calculated power measurement accuracy at each measurement point is a function of the overall current measurement accuracy and the voltage measurement accuracy, as shown in Table 5.

Table 5. Calculated Power Measurement Accuracy

Measurement point	Nominal voltag	Power range	Current measurement accuracy	Voltage measurement accuracy	Calculated power accuracy
Server input					±0.3%
AC system input	480V ac		±1.1%	±0.1%	±1.2%
AC UPS output	208V ac		±1.1%	±0.1%	±1.2%
DC system input	480V ac		±1.1%	±0.1%	±1.2%
DC UPS output	380V dc	5.7 to 68.4kV	±1.1%	±0.1%	±1.2%

These measurement errors need to be considered in determining the difference between the ratio of the measured input powers for a DC and an AC system.

Thus, this analysis of the error propagation in the system revealed that the worst case possible error would be on the order of the energy saved. Therefore, a second set of measurements were carried out (at the same locations) to verify original measurements, using the Power Analyzer and CTs with 0.25% accuracy. Both sets of measurements are included in the results section of this report.

4.6 Other Issues

A number of other issues were identified during our Open House sessions, both technical and non-technical in nature. We list a number of them here:

- Awareness: While the demonstration was able to increase industry awareness of more efficient alternatives to today's practices, the industry's current knowledge of available options is quite low. Additional research and outreach by industry respected sources are needed.

- Creating a Market for DC: For DC approaches to take hold in the data center market, a number of market barriers still need to be addressed in a consistent, unified manner.

- On-Going Efforts: The Demonstration Project accomplished its objective of informing the industry of the DC distribution alternative, but once the project is completed, there are few other places where such a set up can be found.

5. Safety/Codes & Standards/Industry Associations

Due to the experimental nature of the demonstration project, as well as the agreed upon DC voltage level, which is higher than typically found elsewhere, the project team took steps to ensure both equipment and personnel safety.[45] This was accomplished though the establishment of a "Safety Committee" consisting of project participants with extensive experience in data center design and construction, participants from the telecommunication industry, power electronics, and other.

The "Safety Committee" reviewed the implementation plans, one-line drawings and discussed equipment specifications and compatibility prior to assembly. Members of the committee also participated directly in the equipment assembly process. The project team also consulted members of the committee regarding issues such as connectors. Participants also agreed that the Committee was primarily responsible for safety issues "outside of the rack," while equipment manufacturers were responsible for the safety "inside of the rack." This clear division of responsibilities helped to reduce confusion and minimized conflicts during the assembly.

In addition to assisting the project team with equipment finalization and assembly, the Committee also identified issues and relevant standards needed for data center DC power practices. The section below summarizes the recommended practices and associated codes and standards that warrant consideration for a 380VDC power system. At a minimum, these considerations should include:

- Selecting and locating suitable circuit protection
- Grounding and bonding methods to insure proper operation of the suitable circuit protection (as detailed in section 4)
- Leakage currents and shock hazards
- Environmental conditions and maintenance
- Connection Devices
- Power distribution equipment within the rack
- Training and certification of personnel working inside racks
- Training and certification of personnel installing building power distribution
- Installation and maintenance of under floor power distribution
- Warnings and labeling
- Commission Testing

[45] The International Electrotechnical Commission and its national counterparts (IEE, IEEE, VDE, etc.) define high voltage circuits as those with more than 1000 V for alternating current and at least 1500 V for direct current, and distinguish it from low voltage (50–1000 V AC or 120–1500 V DC) and extra low voltage (<50 V AC or <120 V DC) circuits. Note that various safety and insurance organizations consider anything outside of the ELV range (i.e. greater than 50 V) to be dangerous and in need of regulation.

5.1 Selecting and Locating Suitable Circuit Protection

As discussed in section 4, suitable circuit protection using breakers and fuses does not directly translate from AC to DC and suitable DC circuit breakers will likely need to be selected. In terms of the breaker ratings, the manufacturers of these devices must be consulted for coordination. Further, experience in the from the FAA critical DC bus power systems work, along with information from photovoltaic inverter industry, the DC drive industry and electromotive (DC train) industry can serve as a reference resource for further information on this topic. Available fault currents may vary depending upon the type and short circuit current availability from any energy storage devices used. Fault clearing capabilities of DC rated protection devices may not be sufficient under all fault conditions.

5.2 Grounding Methods to insure suitable operation of circuit protection

Section 4 detailed the basic considerations for grounding and bonding, but focuses on grounding topologies and techniques that enhance or supplement basic grounding and bonding requirements. In terms of promoting operation of the circuit protective devices during fault conditions, the recommendations set forth in the National Electric Code (NFPA 780) and the IEEE *Emerald Book* contain the fundamental recommendations that can be incorporated with the previous discussion on suitable circuit protection. Also it should be considered that lower voltages than the subject 380V DC demonstration may be considered in the future, therefore, the telecommunications related standards may provide some of the useful materials related to the lower voltage DC systems.

5.3 Leakage Currents and Shock Hazards

Because AC and DC powered information technology equipment can and does contain noise filters and other components that create leakage currents across the ground paths the requirements outlined in the Underwriters Laboratories document *UL 60950-1* must be taken into consideration as a basic listing for ITE. The document entitled: *Information Technology Equipment – Safety – Part 1: General Requirements* are basically related to leakage current, but the issue of electric shocks is well represented. This relates to considerations where the grounding conductor inadvertently becomes loose or disconnected and the rack or the equipment case becomes energized.

5.4 Connection Devices

Market adoption of higher voltage DC systems will require the standardization of DC connectors that can be safely used by untrained personnel. In existing data centers personnel can connect/disconnect servers powered at 120 V without any safety concerns. Because DC current does not have a zero crossing like AC current, it is more difficult to interrupt, and simple disconnection will cause arcing. In order to address this concern, some connectors have been developed that can contain the arc until it is extinguished. End users will want standard connectors to choose from for these applications.[46]

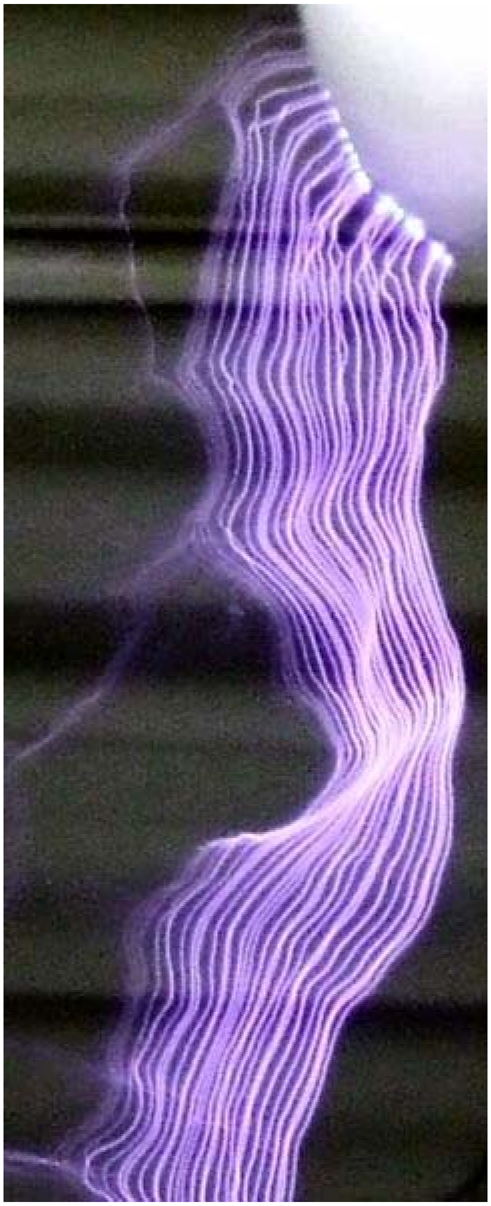

Figure 25. Tesla coil showing repeated electric discharges.

5.5 Power Distribution Equipment within the Rack

For the same reasons outlined in section 5.4 above, cabling and plug strips used in server racks must also have standard safe designs, so that untrained personnel can reasonably be expected to work in the rack without danger.

5.6 Training and Certification of Personnel Working Within Racks

One issue of concern to datacenter users is the availability of trained personnel to work on DC power systems. Licensed electricians are common for AC power systems, but the requirements for handling DC systems will be different and therefore require special training. To address this concern, training programs must be developed and provided to personnel working on DC facilities. This training should include both work safety requirements and recommissioning after any modifications

[46] Interrupting an existing current flow often produces a low voltage spark or arc. As the contacts are separated, a few small points of contact become the last to separate. The current becomes constricted to these small hot spots, causing them to become incandescent, so that they emit electrons (through thermionic emission). Even a small 9V battery can spark noticeably by this mechanism in a darkened room. The ionized air and metal vapor (from the contacts) form plasma, which temporarily bridges the widening gap. If the power supply and load allow sufficient current to flow, a self-sustaining arc may form. Once formed, an arc may be extended to a significant length before breaking the circuit. Attempting to open an inductive load facilitates the formation of an arc since the inductance provides a high voltage pulse whenever the current is interrupted. AC systems make sustained arcing somewhat less likely since the current returns to zero twice per cycle. The arc is extinguished every time the current goes through a zero crossing, and must reignite during the next half cycle in order to maintain the arc.

5.7 Training and Certification of Personnel Working on Distribution Systems

The same requirements are needed for personnel working on facility distribution system as personnel working within equipment racks, discussed above in section 5.6.

5.8 Installation and Maintenance of Under Floor Power Distribution

Power cable under the raised floor does not provide the same level of comfort as overhead cabling, as clearances and touch hazards may be more difficult to control. It is very likely that under floor cabling will be a preferred arrangement, so standards for installing these systems will be necessary for end users to be comfortable with this approach.

5.9 Warnings and Labeling

Because these are not yet common systems, there may be a need to develop supplemental warning labels to place on and inside cabinets where personnel might be exposed to high DC voltages and or capacitive voltages when the system is de-energized

5.10 Commissioning Tests

Once the suitable circuit protection has been designed and integrated into the systems it is important that consistent and repeatable procedures and test methods be developed and made available to personnel who will be performing commissioning of the power systems and the protection settings.

6. Results

The magnitude of the DC efficiency gain is highly dependent on the AC reference system and AC/DC power supply that it is being compared to. In the case of this Demonstration Project, we were fortunate enough to have access to two AC distribution systems as well as two DC conversion/distribution systems, and the efficiency ratios were determined for both sets. Our results indicate that the DC approach does provide an increase in conversion efficiency, however exposing the industry to "best in class" systems may also be useful in improving the average efficiency of data centers.

6.1 AC Reference System Measurements

The AC system efficiency was measured with two different UPS systems. The units were both high-efficiency transformer units loaned by industry partners. The initial measurements and set up primarily were primarily carried out on the original loan unit (AC System A). However, another industry partner was invited to submit their system towards the end of the demonstration period, and complied with their highly efficient AC UPS setup (AC System B).

Measurements here are designated as follows:

- Input energy to UPS= *ACEin*
- Output energy to AC loads= *ACEout*
- Input energy to Rectifier= *DCEin*
- Output energy to DC loads= *DCEout*

The figure below shows how the measurements were made for the AC system:

Figure 26. AC Systems Measurement Set Up

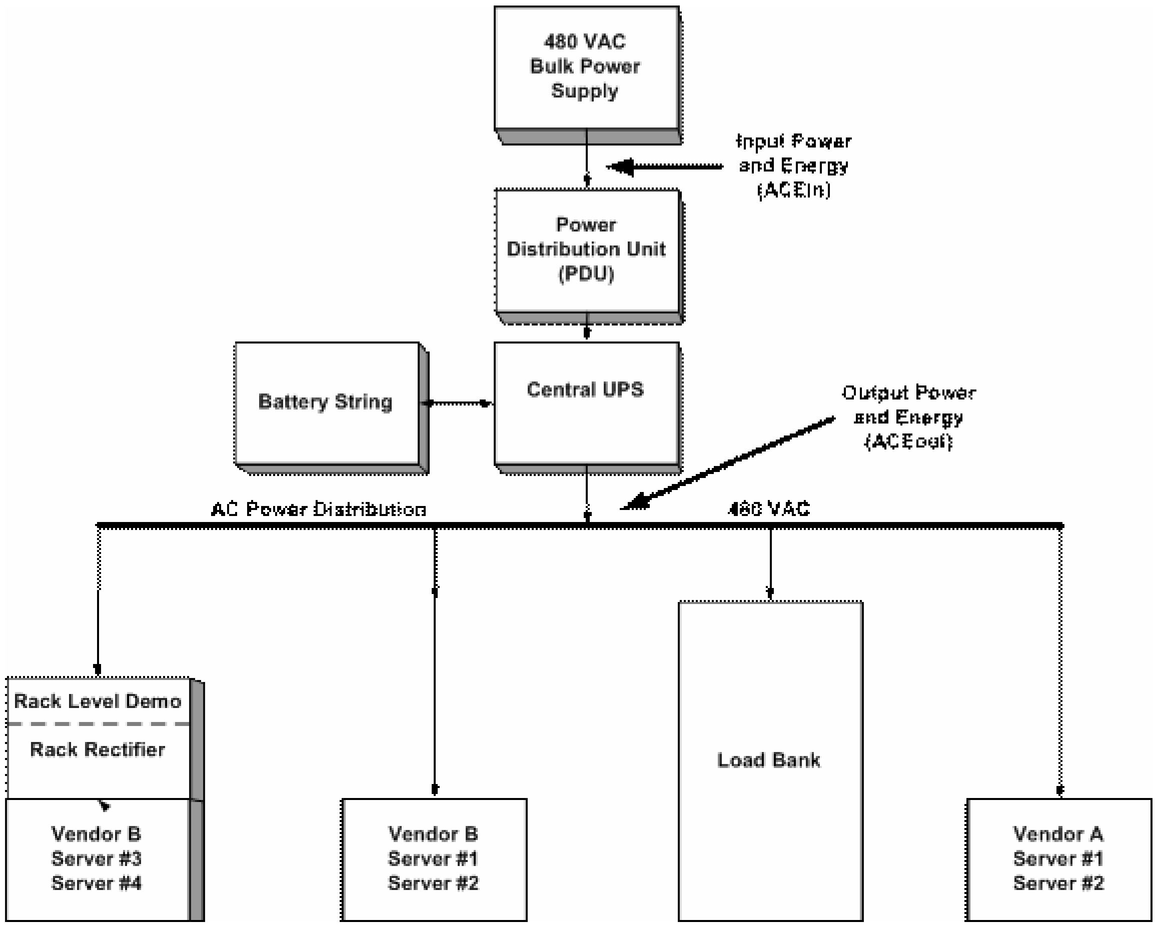

6.2 DC System Measurements

The DC system efficiency was also measured for two different rectifier units, also on loan from industry partners. The efficiency of both rectifier units was measured.

The diagram below shows how the measurements were made for the DC system:

Figure 27. DC Systems Measurement Set Up

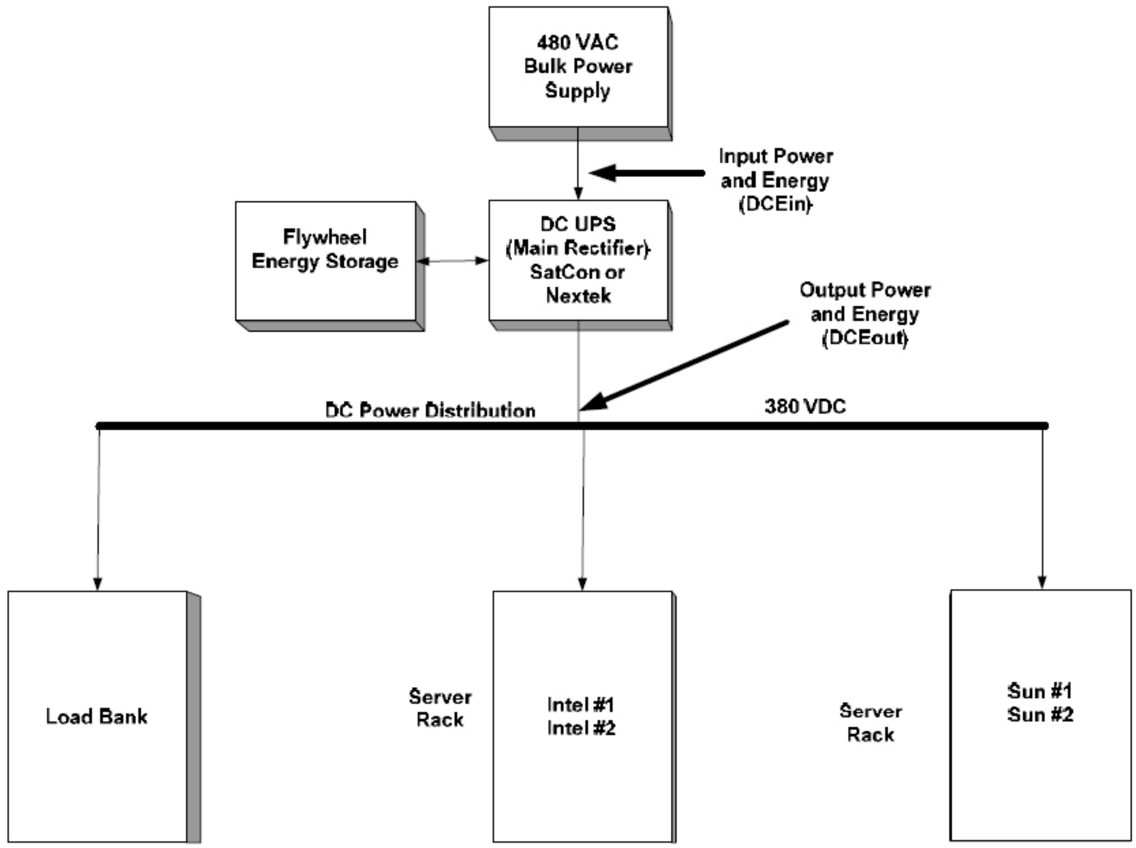

6.3 Efficiencies of AC- and DC-Based Systems

For the measurements, the AC and DC distribution used identical storage and server units except for their power supplies. Both DC and AC processing equipment used the same exercising routines to simulate processor and system loading.[47] Although these "simulated" loads were fairly constant, there were still some small fluctuations. In order to ensure these did not skew the results, power measurements were integrated over typical utility demand periods of 15 minutes, using the following formula:

$$E = \int_0^{15} p(t)dt$$

With the server activities constant for both systems, this means the power on the load side of the AC UPS or DC rectifier is nearly equal, a comparison of the input power can then

[47] Server manufacturers provided their own set up and exercise routines for their equipment sets.

show the difference in energy use. In actuality, there is about a 2% difference in the power loading, due to the higher efficiency of the DC servers.

Therefore,

$$DCEout / ACEout \approx 0.98$$

The objective is to find the percentage improvement, *I*, of the DC system over the AC system.

$$I = \frac{ACEin - DCEin}{ACEin}$$

The AC and DC power measurements are compared using data gathered with DC System A set up first, then using the DC System B set up. Both sets of results are presented below.

Table 6. DC System A Efficiency Measurements

System Efficiency	UPS Efficiency	Transformer Efficiency	PS Efficiency	System Efficiency
AC System A: Measured Efficiency	90%	98%	90%	79%
AC System B: Measured Efficiency	90%	98%	90%	79%
DC System A: Measured Efficiency	94%	100%	92%	87%

Energy Consumption	Compute Load (kWh)	Input Load (kWh)	Efficiency Gain
AC System A: Measured Consumption	23.3	26.0	
AC System B: Measured Consumption	23.3	25.9	
DC System A: Measured Consumption	22.7	24.1	
% Energy Consumption Improvement vs. AC System A			**7.3%**
% Energy Consumption Improvement vs. AC System B			**7.0%**

Table 7. DC System B Efficiency Measurements

System Efficiency	UPS Efficiency	Transformer Efficiency	PS Efficiency	System Efficiency
AC System A: Measured Efficiency	90%	98%	90%	79%
AC System B: Measured Efficiency	90%	98%	90%	79%
DC System B: Measured Efficiency	92%	100%	92%	85%

Energy Consumption	Compute Load (kWh)	Input Load (kWh)	Efficiency Gain
AC System A: Measured Consumption	23.3	26.0	
AC System B: Measured Consumption	23.3	25.9	
DC System B: Measured Consumption	22.7	24.6	
% Energy Consumption Improvement vs. AC System A			**5.0%**
% Energy Consumption Improvement vs. AC System B			**4.7%**

It can be seen in the first table that there is about a 7% decrease in input energy using the DC system compared to the "best in class" AC system. Using the second DC system, the values are slightly lower, but still about 5% improved over the AC system.

Finally, it should be noted that both of the AC distribution system used represent the best on the market with regard to efficiency. Both of the AC UPSs are high efficiency units, and the efficiencies of the power supplies in the AC servers – at 90%, are much higher than most in the data center industry. By comparison, a typical AC system in today's data center would have a UPS that was about 85% efficient, and power supplies around 73% efficient.[48] The estimated improvement of the DC system over these "typical" systems is shown in Table 8. System Efficiency Comparisons below.

Table 8. System Efficiency Comparisons

System Efficiency	UPS Efficiency	Transformer Efficiency	PS Efficiency	System Efficiency
AC Typical Distribution Efficiency	85%	98%	73%	61%
DC Distribution Efficiency	92%	100%	92%	85%

Energy Consumption	Compute Load (W)	Input Load (W)	Efficiency Gain
Typical AC Distribution Efficiency	10,000	16,445	
DC Distribution Option (Optimized)	10,000	11,815	
% Energy Consumption Improvement vs. Typical AC Distribution			**28.2%**

In this case, an improvement of over 28% is possible in an average data center. This means the DC distribution system, as demonstrated, will have the potential of using 28% less energy than the typical AC system found in today's data centers. Since LBNL benchmarking has shown that HVAC loads are typically about the same as the IT load (or worse), this means that a 28% improvement in the efficiency of the electrical distribution to the IT equipment also means a 28% overall facility level improvement.

6.4 Non-Technical Project Results: Open House Sessions

A total of seven Open Houses (including one session on August 18) were held during the duration of the project, with over 200 total attendees representing a diverse cross-section of the data center industry. Attendees ranged from manufacturers to users/operators, government, as well as the media.

A number of Open Houses garnered much larger audiences due to the fact that they were planned to coincide with other events taking place in or near the facility. One was the Data Centre Dynamics conference that took place in San Francisco in July; the other was the Critical Facilities Roundtable quarterly meeting, which was hosted at the facility. The Demonstration also received interest from utilities, and a number of utility representatives and government officials attended a "utility" Open House session.

---

48 These are typical numbers that were found in our evaluation of the servers and UPSs markets.

Figure 28. Demonstration Site with Installed Equipment

The response to the demonstration's press release was quite spectacular, including coverage by a Bay Area television station, as well as articles about the demonstration project by trade and popular press following the press event. The interest generated by the project was quite broad, with coverage by the press ranged from CNN Money to Byte.com, to *CIO Magazine* and *eWeek.* While there was no direct press tracking by LBNL or participants, as of the end of September, the demonstration project generated at over 20 print and on-line articles, and at least one TV interview.[49]

[49] References to the press coverage can be found at http://hightech.lbl.gov/dc-powering/media-coverage.html.

Figure 29. Installed Equipment Map of the Demonstration Site

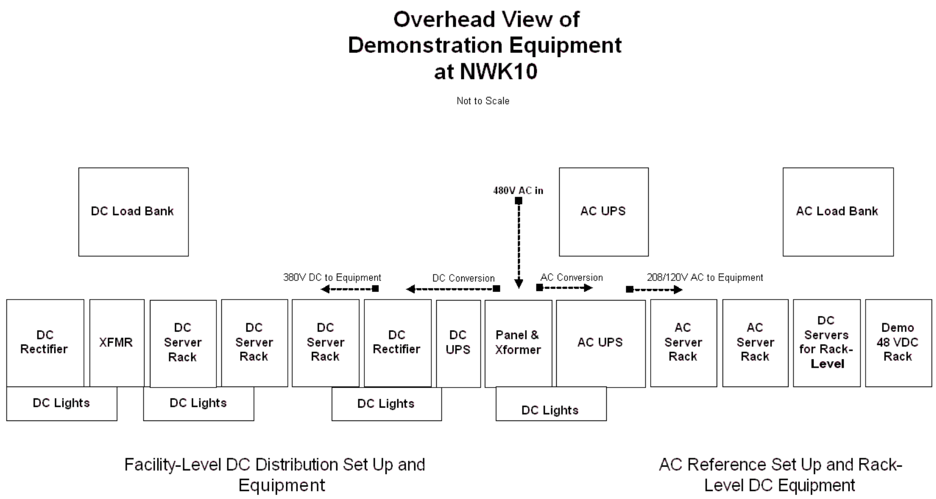

7. Conclusions and Recommendations

In the final analysis, this Demonstration Project was able to coordinate the participation of over 20 companies and their equipment and/or in kind contribution, worked with other companies in the implementation process, and assembled equipment and services worth over a million dollars in value. We were also able to conclusively demonstrate to the data center industry (via the 200+ attendees and the media) that DC delivery systems are viable, can be 20% or more efficient than current AC delivery systems, be more reliable, and potentially cost less in the long run.

Overall, the project succeeded in meeting the objectives that were set out at the beginning. In particular:

> **Availability of DC Equipment**: The Demonstration Project showed that DC-powered servers exist in the same form factor as AC servers or can be built and operated from existing components with minimal effort. The Project reminded the industry that DC servers currently exist (in the 48V DC form factors), and also showed that 380V DC servers and storage equipment could be built and operated from existing components. Further, the Demonstration Project gave visibility to the DC power conversion and distribution equipment, highlighting two commercially available rectification systems, as well as UL-listed buss bars for DC applications.
>
> **DC Functionality**: The Project also showed that DC-powered servers can provide the same level of functionality and computing performance when compared to similarly configured and operating server containing AC power supplies. While the DC-powered servers were not compared with AC powered servers using industry standard software benchmarking tools, this was not a significant issue due to the fact that identical server units (except for the power supplies) were provided by manufacturers. Measurements of efficiency were made using industry standard measurement devices. The Demonstration also included storage units, as well as DC network equipment that can use a variety of DC voltages.
>
> **Demonstrated Gains in Efficiency**: The Project demonstrated efficiency gains from the elimination of multiple conversion steps in the delivery of DC power to the server hardware. Results were measured and documented from two sets of DC delivery systems, and compared to two sets of AC delivery systems. In both cases, the DC delivery system showed a minimum of 5% to 7% efficiency gains without significant optimization over two AC distribution systems that are “best in class” and much more efficient than most systems found in today’s data centers. These measured efficiency gains did not include additional gains from a reduction in cooling loads, which can have the potential for additional savings. Raising awareness of the AC - UPS system efficiency will have a benefit even if the DC solution is not embraced.

Follow-Up Investigations: The Project also identified a number of areas for follow up investigations that will help generate industry discussions, and provide useful leverage points to move the industry forward in the direction of DC distribution. These identified areas are topics that require industry agreement and/or adoption, which included:

- Grounding, Protection and Overloading Prevention: A number of grounding, protection and overload prevention practices for DC data centers are discussed in this report for industry considerations.
- Reliability: Anecdotal data shows that DC-powered data centers have the potential to be more reliable than AC-powered data centers. However, data does not exist for DC voltages higher than 48V DC. Leadership will need to come from the industry in adoption or additional testing for the industry to move forward on this area.
- Costs: While cost data exists, it has not yet been compiled in a way that direct comparisons can be made for the two distribution systems (or their TCO). The Demonstration has generated significant interests from data center designers and system integrators, and further discussion on this area can lead to at least a first-order estimate of DC distribution costs. A related area is the costs of DC components vs. AC components. Currently, AC components may enjoy better economies of scale, but wide-spread DC power adoption may change this equation.
- Integration with Other Sources: A DC data center can also simplify the integration of alternative energy sources, such as solar and other forms of renewable energy, as well as fuel cells and distributed generation, which are all DC-based.
- Measurement Verification: The accuracy of measurements was confirmed with an additional measurement approach. This helped to address any concerns over introduced error factors from the instrumentation used.
- Other Issues: Of significant concerns is the lack of industry knowledge of the advantages of DC distribution, as well as misconceptions about DC power. Additional education and outreach efforts will be required if the energy savings potentials of DC powered data centers are to be realized.

Going forward, there remain many barriers to the adoption of DC power distribution, and need additional follow on work. In particular, a number of barriers have been identified during the course of this projects, they include:

- Increasing Awareness of DC Distribution: The industry’s current knowledge of available options for efficiency and DC distribution is quite low. Further, there is no single, trusted source of information, or an entity dedicated to the promotion of data center energy-efficiency and DC distribution (other than the CEC/PIER efforts). With the current industry support and interest, a “DC Power” association of some sort will help to focus interest and help to elevate awareness among the data center market.

- Creating a Market for DC: A number of market barriers still need to be addressed in a consistent, unified manner. Coordinating utilities and other efforts, at least in California, will go far towards getting DC approaches to take hold in the data center market. There is utility interest in establishing a baseline of performance and cost, which can then help to address at least the early adoption barrier of cost. Other efforts are still needed, and strategies to address market transformation used by the conservation movement can be directly applied here. In addition, the US Congress has recognized the potential for energy savings with HR-5646, so that coordination with DOE and EPA is needed to ensure no duplication of efforts.

- Developing Pilot Projects: Once the Demonstration Project was completed, there are no other places where such a set up can be found. Efforts are needed to continue the Demonstration Project's role in informing the industry of the DC distribution alternative. Discussions are underway with a number of "early adopters," and the project team proposes to establish several pilot projects to:
    - "prime" the market by creating demand for DC servers to enable certification efforts to proceed
    - Determine cost factors for DC systems - capital and operating cost
    - Evaluate and resolve any remaining barriers
    - Publicize successful systems in real data centers

- Develop standards to accelerate adoption: Agreement on distribution voltages, electrical connectors, grounding, DC power quality, and other issues will be important to enable the market to adopt DC distribution on a large scale. The PIER program should facilitate these efforts by bringing together the appropriate industry representatives.

Appendix A: Project Participants

Participants List

Project Organizers

- Lawrence Berkeley National Laboratory: William Tschudi, Evan Mills, Steve Greenberg
- Ecos Consulting: My Ton
- EPRI Solutions: Brian Fortenbery

Industry Partners

- Alindeska Electrical Contractors
- AMD
- APC
- ANCIS Inc.
- Baldwin Technologies
- CCG Facilities
- California Energy Commission
- Cingular Wireless
- Cisco Systems
- Cupertino Electric
- Data Power Design
- Dell
- Dranetz-BMI
- Dupont Fabros
- EDG2
- Emerson Network Power
- EYP Mission Critical Facilities
- Fairchild Semiconductors
- Gannett
- HP
- IBM
- IEM (Industrial Electric Manufacturing)
- Intel
- Morrison Hershfield
- Nextek Power Systems
- NTT Facilities
- Panduit Corp.
- Pentadyne
- Rosendin Electric
- RTKL
- Rackable Systems
- SatCon Power Systems
- SBC Global
- Solara
- Square D
- Sun Microsystems
- 380voltsDC.com
- TDI Power
- Universal Electric
- Verizon Wireless
- Visa

Appendix B: Location for One-Line Drawings of Demonstration Configurations

Demonstration configurations can be found on line at:

http://hightech.lbl.gov/dc-powering/walkthrough.html

Appendix C: Participant "Code of Conduct"

GENERAL UNDERSTANDING AMONG PARTICIPANTS OF THE DC ARCHITECTURE DEMONSTRATION PROJECT

1. This Demonstration will evaluate the energy efficiency benefits of DC power distribution systems in data centers. It will be carried out as a public-private partnership coordinated by Lawrence Berkeley National Laboratory, Ecos Consulting and EPRI Solutions. The work is sponsored by the California Energy Commission and the Participants.

2. There is no cost to become a project Participant. However, it is understood that Participants will contribute equipment, time, and/or expertise.

3. Participants will be identified in project publications and share in the credit and results of the Demonstration.

4. The Demonstration will evaluate data center AC and DC distribution systems from a power delivery efficiency perspective, comparing traditional AC distribution, DC distribution at the rack level, and DC distribution at the facility level. The methodology will be to exercise any given piece of IT equipment identically for both the AC and DC configurations in order to create similar loading on the power delivery systems.
   a. Performance comparison of AC and DC distribution systems will be obtained.
   b. To the extent possible, the vendors of specific equipment will not be publicly identified. Reporting will focus on the energy implications of the distribution systems rather than on the types of equipment or a particular vendor. For example, Rack 1 equipped with servers A, B, and C the configuration may be identified, while the manufacturers of the servers or power conversion equipment will remain anonymous.
   c. Individual equipment (such as a server) may run any application selected by the manufacturer, as long as operating conditions are identical in the AC and DC trials.

5. A final project report will be made available to all participants for review before external submission/dissemination.

6. Participants agree to communicate a consistent message to outside entities regarding the purpose and goals (per this memorandum) as well as the results of this Demonstration (per the final project report).
   a. LBNL will coordinate publicity with Participants and outside entities. Initial news releases should be published on a single day to be mutually determined.
   b. Any "Press coverage" with respect to the Demonstration should clearly identify the goals and results of the project and clearly indicate the "multi-party" involvement in the project.

7. The host facility will make reasonable arrangements for site access.
   a. During setup and operation, access to equipment used in the demonstration should be provided only on an as-needed basis.
   b. No personnel should be able to conduct internal examination of specific equipment other than their own without explicit consent of the relevant equipment owner.
   c. Host facility to provide access (with escort if required) to all participants, their guests, or other interested parties for showcasing the Demonstration.
   d. Organized open house sessions with one of the LBNL team present will be held rather than ad hoc visits.

8. The host facility will provide direct or remote access to monitor power and change of application programs during defined time periods.

Appendix D: Measurement Details

1 AC Reference System Measurements

The AC system efficiency was measured with two different UPS systems. The units were a Mitsubishi-UPS with a standard 75kVA input transformer unit loaned by Baldwin Technologies, Inc. (it is a unit owned by BTI who on loan to Pentadyne for testing with their flywheel system), and an APC-UPS unit with a high efficiency input transformer loaned by American Power Conversion, Inc. The initial measurements and set up primarily used the Mitsubishi UPS system. However, APC was invited to submit their system towards the end of the demonstration period, and complied with a highly efficient AC UPS setup. For thoroughness of the demonstration, the performance of APC UPS was also measured. [50]

2 DC System Measurements

The DC system efficiency was also measured for two different rectifiers. Initial measurements used a SatCon Inc. rectifier unit, but since the Demonstration Project was provided with two DC rectification systems, a Nextek rectifier unit was also measured.

3 Efficiencies of AC- and DC-Based Systems

For the measurements, the AC and DC distribution used identical storage and server units except for their power supplies. Both DC and AC processing equipment used the same exercising routines to simulate processor and system loading.[51] Although these "simulated" loads were fairly constant, there were still some small fluctuations. In order to ensure these did not skew the results, power measurements were integrated over typical utility demand periods of 15 minutes, using the following formula:

$$E = \int_0^{15} p(t)dt$$

With the server activities constant for both systems, this means the power on the load side of the AC UPS or DC rectifier is nearly equal, a comparison of the input power can then show the difference in energy use. In actuality, there is about a 2% difference in the power loading, due to the higher efficiency of the DC servers.

Therefore,

$$DCEout \big/ ACEout \approx 0.98$$

[50] The measurement and monitoring system used was the Signature System on loan from Dranetz-BMI.
[51] Server manufacturers provided their own set up and exercise routines for their equipment sets.

The objective is to find the percentage improvement, I, of the DC system over the AC system.

$$I = \frac{ACEin - DCEin}{ACEin}$$

The AC and DC power measurements are compared using data gathered using the SatCon as the DC rectifier first, then using the Nextek unit as the DC rectifier. Both sets of results are presented below.

Table 9

System Efficiency	UPS Efficiency	Transformer Efficiency	PS Efficiency	System Efficiency
AC System A: Measured Efficiency	90%	98%	90%	79%
AC System B: Measured Efficiency	90%	98%	90%	79%
DC System A: Measured Efficiency	94%	100%	92%	87%

Energy Consumption	Compute Load (kWh)	Input Load (kWh)	Efficiency Gain
AC System A: Measured Consumption	23.3	26.0	
AC System B: Measured Consumption	23.3	25.9	
DC System A: Measured Consumption	22.7	24.1	
% Energy Consumption Improvement vs. AC System A			**7.3%**
% Energy Consumption Improvement vs. AC System B			**7.0%**

Table 10

System Efficiency	UPS Efficiency	Transformer Efficiency	PS Efficiency	System Efficiency
AC System A: Measured Efficiency	90%	98%	90%	79%
AC System B: Measured Efficiency	90%	98%	90%	79%
DC System B: Measured Efficiency	92%	100%	92%	85%

Energy Consumption	Compute Load (kWh)	Input Load (kWh)	Efficiency Gain
AC System A: Measured Consumption	23.3	26.0	
AC System B: Measured Consumption	23.3	25.9	
DC System B: Measured Consumption	22.7	24.6	
% Energy Consumption Improvement vs. AC System A			**5.0%**
% Energy Consumption Improvement vs. AC System B			**4.7%**

It can be seen in the first table that there is about a 7% decrease in input energy using the DC system compared to the "best in class" AC system. Using the Nextek system, the values are slightly lower, but still about 5% improved over the AC system.

Finally, it should be noted that both of the AC distribution system used represent two of the best on the market with regard to efficiency. Both of the AC UPSs are high efficiency units, and the efficiencies of the power supplies in the AC servers – at 90%, are much higher than most in the data center industry. By comparison, a typical AC system in today's data center would have a UPS that was about 85% efficient, and power supplies around 73% efficient.[52] The estimated

[52] These are typical numbers that were found in our evaluation of the servers and UPSs markets.

improvement of the DC system over these "typical" systems is shown in Table 8. System Efficiency Comparisons are shown below.

Table 11

System Efficiency	UPS Efficiency	Transformer Efficiency	PS Efficiency	System Efficiency
AC Typical Distribution Efficiency	85%	98%	73%	61%
DC Distribution Efficiency	92%	100%	92%	85%

Energy Consumption	Compute Load (W)	Input Load (W)	Efficiency Gain
Typical AC Distribution Efficiency	10,000	16,445	
DC Distribution Option (Optimized)	10,000	11,815	
% Energy Consumption Improvement vs. Typical AC Distribution			**28.2%**

In this case, an improvement of over 28% is possible in an average data center. This means the DC distribution system, as demonstrated, will have the potential of using 28% less energy than the typical AC system found in today's data centers. Since LBNL benchmarking has shown that HVAC loads are typically about the same as the IT load (or worse), this means that a 28% improvement in the efficiency of the electrical distribution to the IT equipment also means a 28% overall facility level improvement.