

Evaluating 400V Direct-Current for Data Centers

A case study comparing 400 Vdc with 480-208 Vac power distribution for energy efficiency and other benefits.

**Tomm Aldridge, Annabelle Pratt, Pavan Kumar, Dwight Dupy, Guy ALee
Intel Labs**

Abstract

This paper describes an in-depth case study conducted by Intel® Corporation in conjunction with EYP Mission Critical Facilities (now an HP Company) and Emerson Network Power. The study compares the costs and benefits in selecting a 400 volt direct-current (Vdc) design over the typical North American 480-208 volt alternating-current (Vac) design for a 5.5 megawatt (MW) data center. The study concluded that energy savings of approximately seven to eight-percent could be achieved over high efficiency, best practices 480-208 Vac – with a 15-percent electrical facility capital cost savings, as well as a 33-percent space savings, and 200-percent reliability improvement. The paper points out the differences in configuration between the 400 Vdc and 480-208 Vac designs. It also discusses what is required to enable widespread adoption of 400 Vdc designs.

Introduction

For years, most data center operators focused primarily on providing enough compute capacity to meet demand and provide adequate cooling. But as businesses, organizations and governments have become increasingly dependent on data centers for practically everything they do (from day-to-day operations to electronic transactions and communications), the cost of powering these data centers has gained greater attention. Equally worrisome, many facilities have outgrown the power capacity for which they were designed.

The total power consumption of today's data centers is becoming noticeable. According to the U.S. Environmental Protection Agency (EPA), data centers and servers in the United States accounted for approximately 1.5-percent of the U.S. total electricity consumption in 2006.¹ To put this in perspective, the EPA notes that this is more than the electricity consumed by the nation's color televisions and is “similar to the amount of electricity consumed by approximately 5.8 million average U.S. households (or around five percent of the total U.S. housing stock).” While data centers enable energy savings

in the rest of the economy by allowing us to “work smarter,” e.g. through telecommuting or banking on-line, their growing energy utilization is something the EPA thinks that we can address by adopting energy efficient technologies such as 400 Vdc power distribution.

What’s surprising is that these problems are continuing despite some of the greatest energy efficiency gains per unit of processing power the world has ever seen. In March 2007, for instance, Intel introduced two energy-efficient 50-watt server processors, the Quad-Core Intel® Xeon® processor L5320^A and L5310^A. These products delivered a 35- to nearly 60-percent decrease in power from Intel’s 80- and 120-watt quad-core server products, yet provide similar performance. These 50-watt processors represented a nearly ten-fold improvement in power consumption per core in just 1½ years.²

Intel processors introduced in 2008 offer even greater performance per watt. Compared to 2007’s 65nm Quad-Core Intel® Xeon® processor L5318^A (1.6 GHz), the new 45nm Quad-Core Intel® Xeon® processor L5408^A coupled with the new low-power Intel® 5100 Memory Controller Hub (MCH) chipset delivers up to a 67-percent performance increase per watt.³ That’s 1.67 times more performance per watt in just one year.

Intel processors with continued improvements in performance per watt are planned into the future. At the same time, Intel recognizes that the processor is only one part of the problem and, several years ago, started researching performance per watt with a focus on energy efficient systems architecture at the platform level. When one takes such a holistic, end-to-end view, it is clear that there are significant system opportunities spanning from the utility feed into the enterprise all the way down to the processors, chips, fan(s) and memory. According to American Power Conversion, Inc., a producer of data center power conditioning and UPS gear, much less than half the electrical power feeding a data center is actually delivered to Information Technology (IT) loads.⁴ The rest is lost in power conversion and transformations to supply the equipment and in powering cooling systems. For this reason, many organizations around the world – such as The Green Grid, EPA, European Code of Conduct, Japan’s Ministry of Electronics, and the Technology Information Technology Industry Council (ITI) – are looking into ways to improve the efficiency of these systems and equipment.

Power distribution figures prominently in getting this huge efficiency gain – particularly in the data center. In fact, the power conversion efficiency in a typical data center using 480 Vac power delivery (excluding facility cooling systems) is approximately 50-percent (see figure 1, below). Our research has identified an opportunity to move that to 75-percent efficiency with a systemic approach. Thus, when we look at a typical data center, there is significant opportunity for efficiency gains in the power distribution just between the utility feed and each server (up to 28-percent versus current practice in North America⁵ and five-seven% versus high-efficiency AC⁶).

Intel is a major player in these efforts and is involved in active research to improve data center energy efficiency.⁷ Recently, Intel decided to investigate the benefits of bringing 400 Vdc power directly to the IT equipment. Most typical North American data centers

use 480-208 Vac power delivery (referred to from here on as 480 Vac). In a 480 Vac power distribution system, medium-voltage (MV) from the utility is stepped down with a transformer to 480 Vac at the building entrance or major low-voltage substations, and then stepped down again within the facility to 208/120 Vac (hereinafter as 208 Vac) for distribution within the data center. The goal of this study was to compare the potential costs and benefits in using a 400 Vdc design over a 480 Vac design for a 5.5 MW data center.

The results were encouraging. The study found that using 400 Vdc power delivery would result in an energy savings of approximately seven percent, as well as a 33-percent space savings, a 200-percent reliability improvement, and a 15-percent electrical facility capital cost savings for the facility evaluated. Reliability goes even higher to 1000-percent improved with the 400Vdc bus connected directly to the batteries. Specific improvements may vary for other facilities. This paper discusses the differences in configuration between the 400 Vdc and 480 Vac design and how the results were achieved. It also discusses what would be required to enable widespread adoption of 400 Vdc designs.

Why look for new solutions?

In an age of growing concern for global warming and sustainability, the inefficiencies in data centers, particularly in power delivery, are significant, expensive, and ecologically unnecessary.

According to a report by Jonathan Koomey (Project Scientist, Lawrence Berkeley National Laboratory and Consulting Professor, Stanford University) data centers worldwide consumed 123 billion kWh in 2005 and this total is expected to increase 40 to 76 percent by 2010.⁸ In the U.S., data centers consumed 61 billion kWh (USD 4.5 billion) in 2006 and this total is expected to double by 2011.⁹ Obviously, data centers are a growing slice of the overall energy consumption. What's more, according to Gartner, half of all data centers are already power-constrained and unable to expand.¹⁰ Given that, per the SMART 2020 report, Information and Telecommunications Technology (ICT) has its largest influence "enabling energy efficiency in other sectors" and that ICT "could deliver carbons savings five times larger than the total emissions from the ICT sector in 2020,"¹¹ trying to arrest data center growth would be the exact wrong thing to do. At the same time there is no reason that we have to accept the current level of inefficiency in data center power distribution.

The EPA estimates 10 new power plants will be required by 2011 just to support U.S. data center growth. Powering US data centers in 2011 will account for approximately 70 million metric tons of atmospheric carbon dioxide (CO₂) per year. This is equivalent to about 3 percent of total U.S. CO₂ emissions from electric energy generation in 1999.¹²

Obviously, many things need to be done to improve the efficiency of data centers and reduce their carbon footprint. But where does one start? One particularly promising area for improving energy efficiency is power delivery. In a typical data center, which is not optimized for efficiency, approximately 50 percent of data center power, not including power for cooling equipment, is lost in conversions, transformations and distribution.

Consequently, to find ways to improve efficiency throughout the power conversion chain, Intel decided to study 400 Vdc power distribution.

Intel had six objectives we were hoping to achieve with a different power delivery solution.

- Higher efficiency throughout the power conversion chain
- Lower equipment and installation costs
- Easy integration with the growing number of alternate power sources (solar, wind, fuel cells, etc.)
- Efficiency benefits across all data center electrical systems (lighting, cooling, energy storage), not just compute loads
- Improved reliability and availability
- Space savings through reducing the footprint of power conversion equipment

The third objective, easy integration with alternative power sources, is particularly timely as data center operators begin to look into on-location power sources. Google, for instance, has just finished phase one of a 1.6 MW solar panel installation at its headquarters in Mountain View, California.¹³ Since solar panels and fuel cells generate dc, and wind turbines with variable speed generators use a dc bus to enable connection to the fixed frequency grid, a data center using dc power delivery would enable better utilization of the power generated by these alternate sources.

Optimizing power distribution through a 400 Vdc design

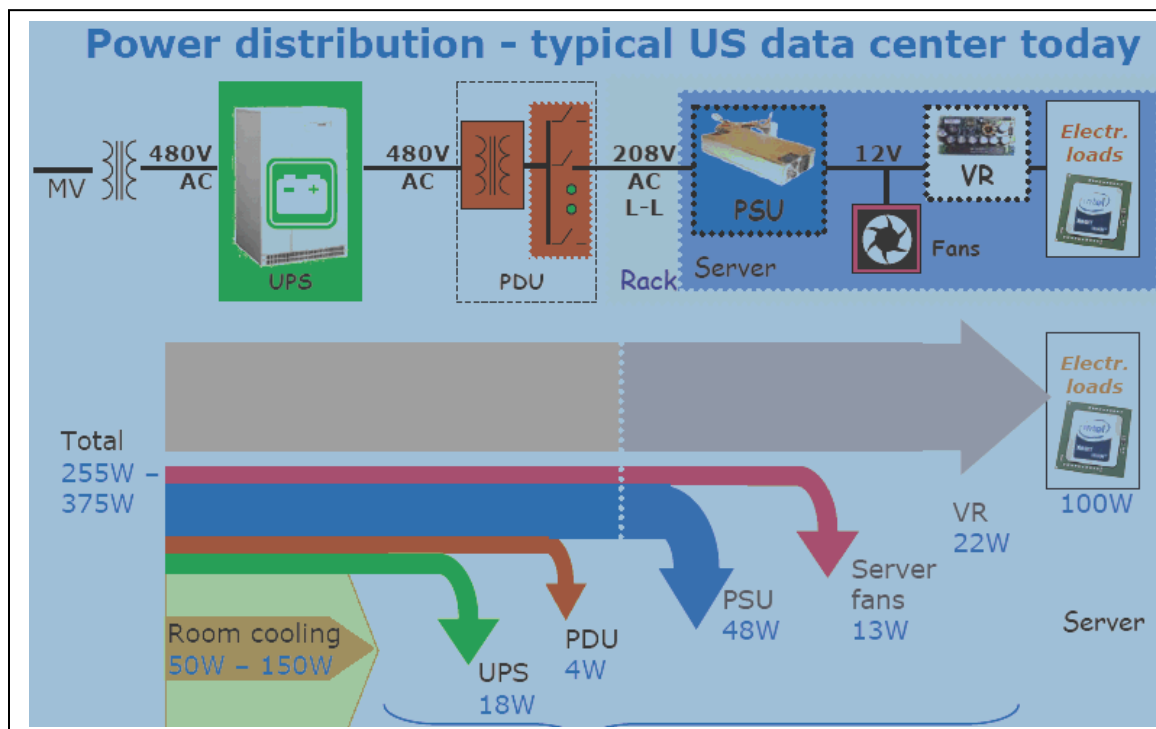


Fig 1. The power distribution efficiency of a typical data center using 480 Vac power delivery (excluding facility cooling systems) is approximately 50 percent.

Figure 1 shows many of the places where power is lost in a typical data center using 480 Vac power delivery. As one can see, in order to deliver 100 watts (W) to the electronic loads, a total of from 255 W to 375 W of power is required at the input to the data center. Anywhere from 50–150 W is used for cooling – an aspect of data center power consumption outside the focus of this paper, but being researched and addressed by Intel and many other companies and organizations.¹⁴

Of the power not being used for cooling, about 50 percent of the power is lost in power distribution. Devices accounting for these losses include:

- Uninterruptible power supplies (UPS)
- Power distribution units (PDU) and cabling
- Power supply units (PSU)
- Voltage regulators (VR)
- Server fans used to cool electronic loads

A 480 Vac data center using best-in-class components can improve the power conversion and distribution efficiency up to approximately 68 percent. This is an area in which many data center operators are currently looking for solutions.

Intel, in collaboration with various industry partners, is taking a systems level approach by looking at alternative power delivery architecture. 400 Vdc power delivery is particularly promising because it eliminates three power conversion steps and enables single end-to-end voltage throughout the data center. Previous studies have also identified 400 Vdc power distribution as the most efficient, including the Intel technical paper, “Evaluation of 400 Vdc Distribution in Telco and Data Centers to Improve Energy Efficiency.”¹⁵ Different nominal voltages have been proposed, e.g. 380 Vdc was used in the small scale test model assembled in collaboration with Lawrence Berkeley National Laboratory (LBNL) and other industry partners (a full description of the setup can be found in the report issued by LBNL¹⁶).

While Intel recommends 400Vdc as the nominal voltage and used 400Vdc for this study, Intel supports international harmonization of the voltage level and looks to standardization efforts to finalize the specification.

Additional advantages of 400 Vdc power distribution include the number of ways it simplifies power management.

- No phase balancing is required, which reduces the complexity of power strips and wiring.
- No synchronization is required to parallel multiple sources
- There are no harmonic currents to worry about, eliminating the need for power factor correction (PFC) circuits.
- It can use fewer breakers (up to 50 percent for this case study) because of fewer power conversion stages.
- It simplifies wiring since only two conductors are required.

The use of 400 Vdc is considered preferable over a higher or lower voltage for a variety of reasons.¹⁷

- A link voltage of approximately 400Vdc already exists in today's ac power supplies, as well as the bus in light ballasts and adjustable speed drives (ASDs) which are often used to power fans and pumps in the data center. Because 400 Vdc builds on the existing components in high volume production server and desktop power supplies, the lowest power supply cost is assured by not increasing the voltage any higher.
- UPS systems typically use a higher voltage dc bus of 540Vdc which can easily be re-designed to support 400 Vdc.
- The spacing requirements per the IEC 60905-1 standard for power supplies are the same for universal input (90 – 264 Vac) power supplies and for dc power supplies with a working voltage below 420 Vdc. This is advantageous as it would allow the re-use of high-volume designs in existence for ac power supplies.
- It is well within an existing 600 V safety limit.
- It operates over standard 600 V rated wiring and busing systems.
- Commercial solutions are already emerging.
- It is simpler and more efficient to connect to renewable energy sources such as photovoltaics, fuel cells, and wind turbines with variable frequency generators, since they also already have a higher voltage dc bus at a voltage in this range.

Comparing dc versus ac for a planned addition to a data center complex

To determine the benefits of installing a 400 Vdc data center, Intel, in collaboration with Emerson Network Power and EYP Mission Critical Facilities, analyzed an expansion module project to an existing data center complex.¹⁸ Each module in the complex houses 220 racks of IT equipment in a two-story, vertical flow-through air design with a centralized cooling plant. The facility was originally designed to contain five 3.3 MW modules at 15 kilowatts (kW) per rack and two of these have been built-out and are in operation. Due to the increase in rack power density, the third module needs to be designed for 25 kW/racks for a total 5.5 MW IT load. A design for this module had been completed based on conventional 480 Vac distribution. The case study asked the question: How would one do this addition with a 400 Vdc design and how would it compare to the planned 480 Vac design? The study evaluated the potential benefits of this 400 Vdc design.

Figure 2 compares the power distribution path for both the studied power delivery systems, 480 Vac and 400 Vdc. It can be seen, that both start with transforming the medium-voltage supply from the utility to 480 Vac power to the building. To optimize efficiency, the rectifier is proposed to be a non-isolated topology with the medium-to-low voltage transformer providing the required isolation.

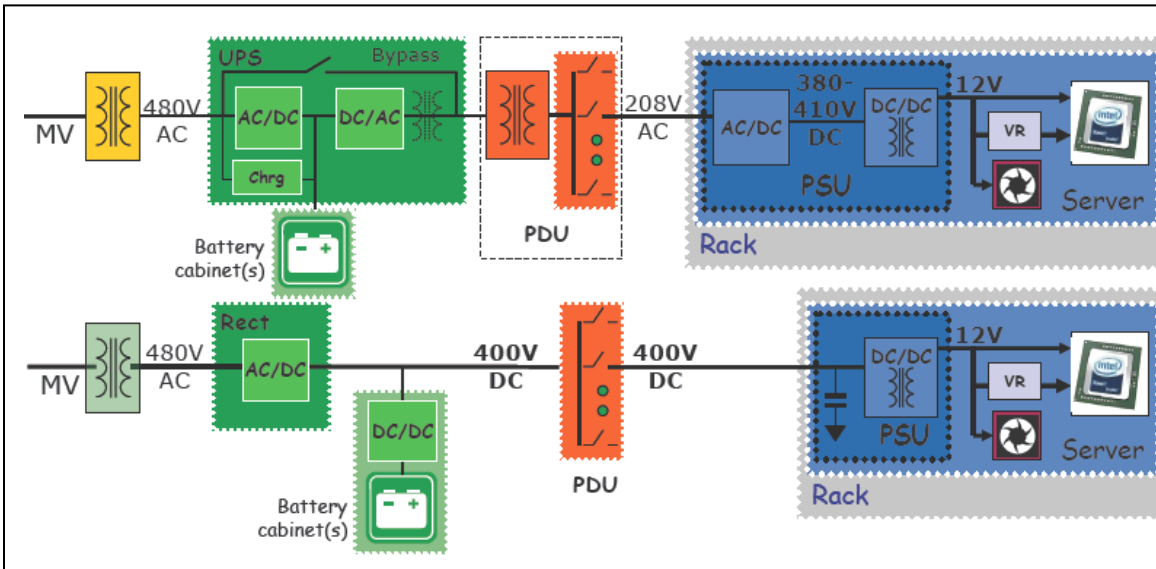


Fig 2. Comparison of the power distribution path for both power delivery systems, 480 Vac (top) and 400 Vdc (bottom).

In the ac delivery system, the power immediately undergoes a double conversion in a centralized UPS which supplies power to many racks. Power is rectified to dc to feed an intermediate backup storage system, and then it is inverted back to ac and sent to the PDU. At the PDU, voltage is stepped down to 208/120 volts to feed each server in the rack. The PSU in each server rectifies the 208 Vdc to a dc voltage typically between 380 and 410 Vdc. The 380 volts dc is then converted with a dc-to-dc converter to 12 volts dc. Some loads, such as hard drives, can take 12 volts directly. Other loads, such as processors, need VRs to step down the voltage.

In the dc power delivery system, one can see the power path is much more streamlined. It removes a dc-to-ac inversion stage in the UPS and an ac-to-dc rectification stage in the PSU. There is also no need for a transformer in the PDU.

It is instructive to look at a PSU to see how dc power delivery can streamline the components themselves (See Figure 3). The removal of the ac-to-dc conversion stage, which provides power factor correction (PFC), means approximately fifteen-percent fewer components by volume and a corresponding reduction in heat generation in the rack of approximately 1.5 percent.¹⁹ Removing these




Fig 3. Switching to 400 Vdc power distribution results in approximately 15 percent fewer components by volume for PSUs. Board picture (prior to callouts) courtesy of Delta Electronics, used with permission.

components saves materials and could reduce the size of the PSU, which frees up additional space on the rack for increased computing density.

The amount of e-waste we have to deal with is staggering. It is estimated that 500 million computers will have been discarded in the US by 2007, resulting in 6.32 billion pounds of plastic and 1.58 billion pounds of lead.²⁰ Therefore, minimizing the number of components required to provide the desired functionality is becoming an ever more important consideration.

In this proposal, a dc/dc converter was included between the battery and 400 Vdc distribution bus, as this provides a regulated output. It is also possible to connect the batteries directly to the 400 Vdc distribution bus, similar to what is done in -48 Vdc telecommunication facilities. This would impose a wider input voltage range requirement on the dc/dc converter in the PSU than what is typical in an ac power supply, which would remove the benefit of re-using high-voltage designs from ac power supplies. However, modern data centers can bring back-up diesel-fueled engine-generators on-line in less than a minute, during which time the battery voltage should stay within the PSU's dc/dc converter input range. Such a direct battery connection would provide further reliability and cost improvements.

Figure 2 shows a simple one-line diagram of the power distribution. Because of the high availability requirement, a certain level of redundancy needs to be built into the distribution, making the distribution more complex. For the 480 Vac facility design, an isolated-redundant (iso-redundant) configuration with two sets of five 800 kVA UPS systems (where one of the five UPSs in each set provides iso-redundancy) was used. This is sometimes referred to as a 2(N+1) configuration, where "N" is the required number of UPS units needed to condition the data center load being served.

A distributed-redundant configuration was selected for the 400 Vdc facility to optimize its efficiency.²¹ Furthermore, each rectifier, as shown in Figure 2, was proposed to be made up of three 550 kW rectifier modules, with one of them being redundant. The rating of each block is therefore 1.1 MW, requiring five rectifiers for the facility.

Estimated total energy savings

To figure out the total potential energy savings of 400 Vdc versus 480 Vac power delivery for the 5.5 MW data center addition, Intel used an analytical modeling tool they developed to compare the efficiency of different power distribution architectures. More details can be found in "Evaluation of 400 Vdc distribution in telco and data centers to improve energy efficiency," *IEEE Intelec 2007*²², where a general comparison is provided. The tool is based on extracting analytical expressions for the losses in each conversion stage from measured data. For this case study, the tool was used with data from specific equipment from the ac and dc designs.

The tool evaluates power delivery efficiency over a range of loads since converters in the power delivery train often operate at less than half of their rated load and where conversion efficiency is much lower than at heavy loads. This is because many server

workloads include significant idle time, power conversion equipment is often de-rated, and high availability requirements drive the use of redundant power delivery architectures.

The two distribution systems are modeled appropriately for each design, i.e., an iso-redundant configuration for the 480 Vac facility and a distributed-redundant configuration for the 400 Vdc facility. A 2.15-percent total power loss from cabling for both setups is modeled based on typical allowed voltage drops. Modeling of both non-redundant and redundant PSU cases were done.




Fig 4. This graph shows the estimated total energy savings for all the designs modeled. Naturally, redundant power supplies diminish efficiency, but note that the overall advantage of 400 Vdc power distribution remains about the same in the systems using redundant PSUs.

The results were telling, as seen in Figure 4, above. It was estimated that the proposed 400 Vdc power distribution would deliver approximately a seven-percent energy savings over the 480 Vac design. (Actual figures were 7.7-percent savings at 50-percent load and 6.9-percent savings at 80-percent load.) The main source of these savings was the elimination of three power conversion steps. The contributions in estimated savings from each of these stages at the 50-percent load point are as follows:

- The use of high efficiency PSUs was assumed; i.e., with heavy load efficiency of the ac PSU of approximately 88 percent. Based on measurements of these ac PSUs, modified to run off 400 Vdc, a savings of ~1.5 percent was attributed to the removal of the PFC stage.
- The PDU transformer contributed ~1.5 percent, based on published transformer efficiency data.

- Based on detailed simulations and experimental results from a small scale prototype developed by Emerson Network Power, the rectifier efficiency was estimated to be ~97-percent efficient, whereas the 480 Vac design specified a 94-percent efficient UPS, resulting in ~3 percent savings.

The significance of these energy savings becomes more apparent when computed into monetary savings. For a 5.5 MW data center operating at 50-percent load and assuming the average US electricity cost of USD 6.2 cents per kWh²³, 400 Vdc power distribution would deliver annual savings of USD 150,000. When one considers the earlier statistic that U.S. data centers consumed 61 billion kWh (USD 4.5 billion) in 2006 and this total is expected to double by 2011, that seven-percent savings could reduce that gain by 4.27 billion kWh and save USD 315 million.

Relative capital cost comparison

While lifetime operating costs are where the most significant savings are realized, a 400 Vdc power distribution system also provides significant savings in capital costs. A capital cost analysis shows that a 15-percent reduction in electrical equipment (see Table 1) can be achieved by using 400 Vdc power distribution. Since the electrical equipment costs are about 40 percent of a facility’s total costs, this can mean a six-percent savings on the total cost of building a facility. In order to make a direct cost comparison to the cost analysis based on the existing 480 Vac modules, an iso-redundant 3.3 MW version of the 400 Vdc design was used (assuming that 400 Vdc equipment is produced in similar volumes to ac equipment).

Description	ac Data Center [pu]	dc Data Center [pu]	Ratio	Comments
Labor	0.193	0.156	0.8	Savings results from reduced labor cost since less equipment needs to be installed in a 400 Vdc data center. (No PDUs, fewer circuit breakers, etc.)
Materials	0.215	0.227	1.06	Includes wires, conduits, supports, panelboards, busways (including circuit breakers, buses, etc). Increased cost in 400 Vdc is due to different type of equipment albeit reduced in number.
Owner-Furnished Items	0.592	0.467	0.79	Includes UPS, static switches, low-voltage distribution transformers (required in ac data center only), server power supply, etc. Equipment supplied by user (Intel).
Electrical Cost of Work (COW)	1.0	0.85		Total cost to implement all electrical work including equipment.

Table 1. Due to smaller wiring sizes, no distribution transformers, and other economies, Intel estimates a 15-percent savings in the electrical implementation costs of a 400 Vdc power distribution system.²⁴ [In this table, the initial electrical capital cost of the ac data center design is given a value of one and all other figures are given as a fraction (in decimals) of that unit.]

Table 2 shows a breakdown of the cost associated with the above mentioned owner-furnished items (see Figure 5). The simpler implementation of 400 Vdc power distribution – with its reduced components, lack of phasing requirements and harmonic-current mitigation, and easier wiring – make a strong argument for considering it on a

capital cost savings basis alone. The use of fewer components also means additional benefits:

- Using less of the earth’s resources – a significant consideration when you consider the sheer number of data centers and that this number is expected to continue growing dramatically.
- Increased availability of the data center since fewer components and connections will result in fewer failures and improved reliability.

Description	ac Data Center [pu]	dc Data Center [pu]	Comments
UPS / Rectifier	1.0	0.83	Inclusive of batteries, dc/dc converter, input and output switchgear
480 / 208 Transformers	1.0	0.0	Distribution transformers are eliminated in a 400 Vdc data center
Server Power Supply	1.0	1.0	Assumed same cost as an ac power supply although a 6-percent reduction in component cost is estimated

Table 2. Per unit cost breakdown of Owner-furnished items listed in Figure 5.

Reliability benefits

Reliability is an important consideration in any switch to a different power distribution system. After all, today’s businesses, organizations and governments depend on data centers operating 24/7. This makes the fewer conversions and components of the 400 Vdc power distribution particularly attractive. Fewer components mean fewer things that can go wrong. And analysis supports this.

Relx* software reliability prediction calculations done by EYP Mission Critical Facilities predict a 200-percent improvement for a two times lower probability of failure over five years compared to an equivalent Tier IV ac facility. [A Tier IV facility is defined as “providing multiple active power and cooling distribution paths, redundant components, fault tolerant, and 99.995-percent availability.”] Others have calculated a 1000-percent reliability improvement could be achieved with a direct connect to batteries.²⁵ Since availability is a key element in any data center operation, these results make 400 Vdc power distribution appealing for the reliability benefits alone.

<u>Option</u>	<u>Availability</u>	<u>Unavailability</u>	<u>Probability of failure in 5 years</u>
AC Tier IV configuration	0.999996	3.9 e -06	13.63%
DC configuration**	0.999998	2.4 e -06	6.72%
DC Improvement		62.5%	200%

Table 3. Reliability analysis results from Relx* software calculations by EYP Mission Critical Facilities.²⁶

Space savings

Another key advantage for using 400 Vdc power distribution in a data center is space savings. For the 480 Vac design, previously discussed in the two-story model, the subfloor footprint (area below the IT equipment) would have to increase by 50 percent in order to support a 5.5 MW IT load in the upper floor, as shown in Figure 5.

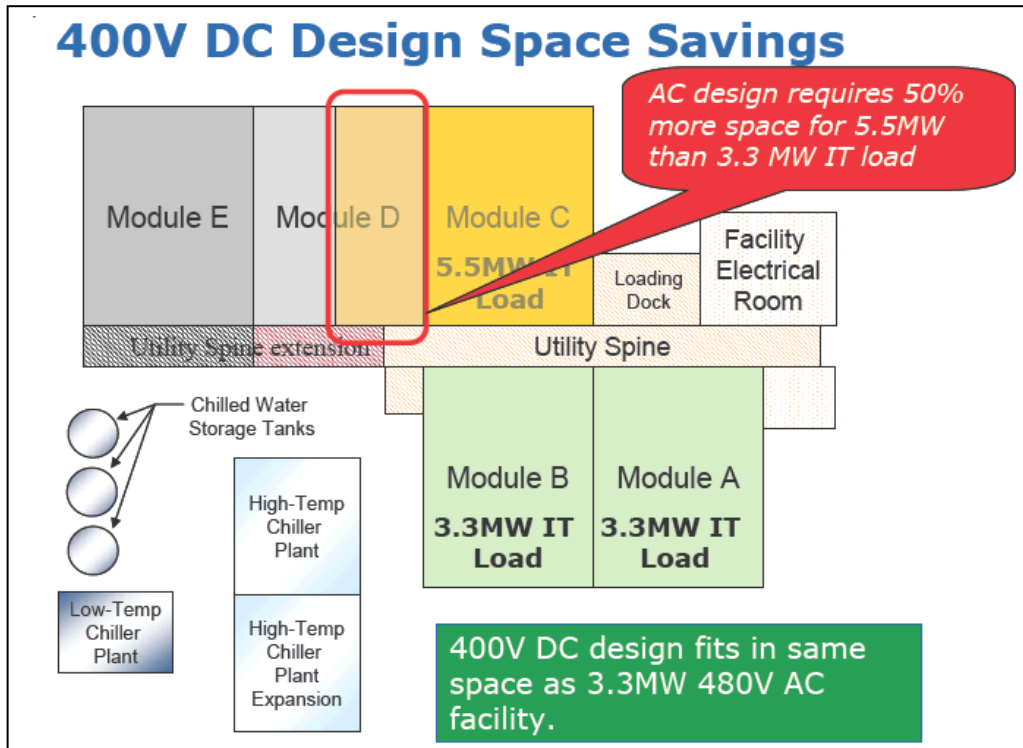


Fig 5. Installing a 400 Vdc design in Module C would enable a 5.5 MW IT load. A 480 Vac design capable of handling the same load would have required 50 percent more space, reducing the amount of space available for a planned Module D.

Fewer power conversion steps and components (no PDUs and simplified switchgear, such as 50-percent fewer circuit breakers) allow a 5.5 MW 400 Vdc design to fit in the same space as a 3.3 MW, 480 Vac design, resulting in a 33-percent space savings over the 480 Vac design. Thus, a 400 Vdc design is a good alternative for packing more compute capacity in the same amount of space.

Making 400 Vdc commercially available

Making the savings of 400 Vdc power distribution available to data center operators will be surprisingly easy. Much of the necessary equipment exists and is in use in other industries. The lack of familiarity with dc power distribution, particularly in the server industry, can be dealt with through education and consultation.

For volume market penetration of this technology there will need to be internationally harmonized specifications for voltage and safe application. While the major safety codes address this dc voltage level, there are differences in local and national building and fire

codes which will have to be harmonized. Intel has proposed a standard for the 400 Vdc distribution bus to the international power community and is promoting discussions of it through The Green Grid technical working group.

Adapting servers to 400 Vdc power requires only changes in the power supply. It entails removal of the PFC stage, keeping the existing dc/dc stage and changing the electromagnetic interference (EMI) filter. PSUs will also need 400 Vdc appliance connectors, but prototypes are already under development and will have Underwriters Laboratories (UL), Canadian Standards Association (CSA) and Technischer Überwachungs-Verein (TÜV) approvals. Other components, such as 400 Vdc rack power strips will also need to be produced. Several companies are working on the development of commercial 400 Vdc rectifiers, including a major supplier of ac UPS systems. None of the necessary components present serious obstacles to implementation of 400 Vdc power distribution. All can become quickly available in response to demand. In fact, items like the dc rack power strips will actually be simpler in design than their ac cousins because they will not need to include features for phase balancing. Many dc quick-connect devices are commonly used today when connecting the dc battery strings to the ac UPS systems. Other items, such as IEC 309 connectors used for dc connections, are certified in Europe and simply need to achieve certification in North America.

As far as maintenance is concerned, all dc couplers and connections can be designed for the same safe access and plug in/disconnect procedures with which users of ac equipment are accustomed. Appliance couplers would have UL, CSA and TÜV approvals. Rack connection and disconnect would be dc-rated IEC 309 or Anderson Power Products® connectors. Racks would be rated for hot disconnect, although we recommend the use of interlocks. Rectifier servicing would be the same as ac UPS and addressed by the ability to isolate each rectifier.

One common concern voiced by data center owners and operators is the safety of personnel when a distribution voltage of 400 Vdc is present. Current data centers with conventional UPS system, which contain batteries as their backup energy source, are operating with dc bus voltages at the UPS at about 540 Vdc. Operation and maintenance personnel are presently familiar with voltages at these levels, and the use of a 400 Vdc system is not more hazardous than what is presently used. It is known that the use of higher ac or dc voltages inside of the data center racks and cabinets may expose data center personnel to voltages with which they are not normally familiar. In the U.S., however, any nominal voltage above 50 volts is considered hazardous to personnel, and when one is using or working with a higher voltage than this, the same precautions and personal protective equipment are mandatory.

Summary

400 Vdc provides the highest efficiency, cost effective solution to powering the Data Center. Intel's case study of a proposed 400 Vdc power distribution compared to the planned design based on conventional 480 Vac distribution for a potential addition of a 5.5 MW expansion module to an existing data center reveals improved energy efficiency over a broad load range. The seven to eight-percent energy savings alone provide reason

enough to consider the switch. When one considers the additional benefits of 400 Vdc power distribution and the ease of integration with alternative energy generation, the argument becomes even more compelling. The fifteen percent less capital cost, 200-percent reliability improvement, and 33-percent facility space savings, make a strong case for the use of dc power distribution to not just improve the energy efficiency of future data centers, but also to reduce the costs of building them and improve the overall availability of their systems. What's more, the transition to 400 Vdc power distribution would be rather easy, requiring just the manufacture and certification of a few simple components that are easily derived from current products. In fact the certification is no different from what would be required to certify costlier high-efficiency Vac alternatives. Given the imperatives to save cost, save energy and reduce carbon emissions, 400Vdc is not only compelling, it appears to be inevitable.

Learn more

To learn more about DC power for data centers, see <http://hightech.lbl.gov/dc-powering/>

To find out more about Intel's efforts in improving data center efficiency, see: <http://www.intel.com/technology/eep/data-center-efficiency/index.htm>

For specific questions or feedback related to this whitepaper please contact the authors by email to guy.allee@intel.com.

* Other names and brands may be claimed as the property of others.

Δ Intel® processor numbers are not a measure of performance. Processor numbers differentiate features within each processor series, not across different processor sequences. See http://www.intel.com/products/processor_number for details.

¹ *Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431*, US EPA, ENERGY STAR Program, August 2, 2007, pp. 10-11.

² Compared to 64bit Intel® Xeon® Processors at 110W TDP (110W per core).

³ Performance comparison of 2x Quad-Core Intel Xeon processor L5318 on an Intel 5000P Chipset-based platform versus 2x Quad-Core Intel Xeon processor L5408 on an Intel 5100 Chipset-based platform using the same workloads. Actual performance may vary. Source: SPECint*_rate_base2006* (score) benchmarks are Intel internal measured results as of January 2008.

Configurations:

- Configuration #1: Processor: Quad-Core Intel® Xeon® processor L5318 (Clovertown), quad core, 1.6 GHz, 8MByte L2 Cache, 1067 MHz FSB. Platform: Intel® Server Board S5000PSLSATAR (Starlake) Reference Board, Intel 5000P Chipset (Blackford), 4x2GB FBD-667 Dual Rank. Software: SuSE 10.1 RC 5 32bit Linux, 2.6.16.46-0.12-smp kernel, CPU2006 V1.0 benchmark suite, Intel ICC10.1 compiler, American Megatrends Inc BIOS 84.
- Configuration #2: Processor: Quad-Core Intel® Xeon® processor L5408 (Harpertown), quad core, 2.13 GHz, 12MByte L2 Cache, 1067 MHz FSB. Platform: Intel® Williamsport Customer Reference Board, Intel 5100 Chipset (San Clemente), 4x2GB DDR2-667 Dual Rank. Software: Red Hat Enterprise Linux 4 Update 2, 2.6.9-22.ELsmp kernel, CPU2006 V1.0 benchmark suite, Intel ICC10.1 compiler, American Megatrends Inc BIOS 80014.
- Benchmark Description for SPECcpu*2006 suite (SPECint*_rate_base2006): SPEC CPU2006 is the industry adopted, CPU-intensive benchmark which stresses the system processor(s), memory subsystem, and compiler. Derived from 29 real user applications, CPU2006 provides a comparison across the widest practical range of hardware reporting a geometric mean ratio score on a baseline compiled binary.

- SPEC benchmark tests results for Intel® microprocessors are determined using particular, well-configured systems. These results may or may not reflect the relative performance of Intel microprocessors in systems with different hardware or software designs or configurations (including compilers). Buyers should consult other sources of information, including system benchmarks, to evaluate the performance of systems they are considering purchasing. Additional information for CPU2006 can be found at www.spec.org. Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering for purchase. For more information on performance tests and on the performance of Intel products, visit http://www.intel.com/performance/resources/benchmark_limitations.htm.

4 Rasmussen, Neil, *Electrical Efficiency Modeling of Data Centers*, American Power Conversion, 2005.

5 Tschudi, et.al., “DC Power for Improved Data Center Efficiency”, Lawrence Berkley National Labs, January 2007.

6 The Green Grid, “Peer Review of Lawrence Berkeley National Laboratory (LBNL) Study on Direct Current in the Data Center”, May 2008. http://www.thegreengrid.org/gg_content/White_Paper_12_-_LBNL_Peer_Review05.09.08.pdf

7 For more on Intel’s current work in data center efficiency, see <http://www.intel.com/technology/eep/data-center-efficiency/index.htm>

8 Koomey, Jonathan G., *Estimating Regional Power Consumption by Servers*, 2007.

9 *EPA Report to Congress on Server and Data Center Energy Efficiency*, August 2, 2007.

10 *Gartner Says 50 Percent of Data Centers Will Have Insufficient Power and Cooling Capacity by 2008*, Gartner Inc. press release, November 29, 2006.

11 *SMART 2020: Enabling the low carbon econoby in the information age*, The Climate Group on behalf of the Global e-Sustainability Initiative (GeSI), ©2008, p. 10.

12 *Carbon Dioxide Emissions from the Generation of Electric Power in the United States*, Department of Energy and Environmental Protection Agency, July 2000.

13 For more on Google’s energy achievements, see: *The Greening of Google* By Sandra Upson, IEEE Spectrum, October 2007.

14 See these Intel white papers: *The State of Data Center Cooling* (<http://download.intel.com/technology/eep/data-center-efficiency/state-of-date-center-cooling.pdf>) and *Air-Cooled High-Performance Data Centers: Case Studies and Best Methods* (<http://www.intel.com/it/pdf/air-cooled-data-centers.pdf>)

15 Aldridge, Tomm V., Pavan Kumar, and Annabelle Pratt, *Evaluation of 400V DC Distribution in Telco and Data Centers to Improve Energy Efficiency*, 2007. Also see: Rasmussen, Neil, *AC vs. DC Distribution for Data Centers*, version 5, 2008 (http://www.apcmedia.com/salestools/SADE-5TNRLG_R5_EN.pdf)

16 http://hightech.lbl.gov/documents/DATA_CENTERS/DCDemoFinalReport.pdf

17 For more detail, see above source (Aldridge) and Kumar, Pavan and Annabelle Pratt, *Evaluation of Direct Current Distribution in Data Centers to Improve Energy Efficiency*, The Data Center Journal, March 2007 (http://hightech.lbl.gov/documents/DATA_CENTERS/DC-Journal-March07.pdf)

18 For a virtual tour, see: http://intelpr.feedroom.com/?skin=oneclip&fr_story=4e2f211fbd4eb8c48e31b558971ea3ba85089725&rf=ev&autoplay=true

19 Rasmussen, Neil, *AC vs. DC Distribution for Data Centers*, version 5, 2008.

20 Gawel, *Groups Aim to make molehill out of mountain of electronic trash*, February 19, 2001 (<http://electronicdesign.com/Articles/ArticleID/4185/4185.html>).

21 For more on distributed redundant configurations such as this, see A. Pines, *Advantages of MFR (Multiple Feed Redundancy) for data centers*, INTELEC 2007, Sept. 30 2007-Oct. 4 2007, Rome, Italy, pages 666 – 668.

22 Aldridge, Tomm V., Pavan Kumar, and Annabelle Pratt.

-
- ²³ Average price of U.S. electricity in November 2007 when this study was performed. Source : http://www.eia.doe.gov/mer/pdf/pages/sec9_14.pdf Energy Information Administration, *Monthly Energy Review*, August 2008, Table 9.9.
- ²⁴ Cost comparison assumptions include:
- Cost comparison analysis performed for 3.3MW IT load
 - Electrical cost of proposed ac design based on existing Intel facility (Module A)
 - Cost comparison for maturity, i.e., high volume
 - Both ac and dc UPSs are iso-redundant
 - Where commercial equipment is not available:
 - Rectifier: two sources
 - Actual cost estimates from Emerson Network Power for 1.1MW (N+1) rectifier
 - ac to dc UPS ratio from Emerson Network Power based on 550kW rectifier module
 - CDU: assume similar cost
 - IEC 309 standard international connector
 - Remove phase balancing complexity
 - PSU: fewer components due to removal of PFC
 - 1st pass estimate of similar cost in comparable volumes
 - Emerson Network Power: BOM cost reduction of ~six-percent
 - Delta: ~1.5x at 20-percent of ac volume (3 yrs), 1.1x at same volume (5 yrs). (Other power supply manufacturers have indicated dc PSU likely to be same cost net of fewer components, but lower volume. Assumed higher cost for this analysis.)
 - 480Vac Module A used as a baseline
 - 2x 15kV -> 480V transformer upstream of UPS
 - 3-phase 480Vac input to UPS
 - 5x 800kVA/720kW UPS with iso-redundancy
 - 3-phase, 5 wire distribution (3 phases, 1 neutral, 1 ground)
 - 480V/208V transformers used in PDUs
 - Proposed 400Vdc design assumptions
 - 2x 15kV -> 480V transformer upstream of UPS
 - No change in the distribution and cost upstream of UPS
 - 3-phase 480Vac input to UPS
 - 5x 720kW dc UPS (rectifier + dc/dc converter in the battery cabinet) with iso-redundancy
 - 5 wire distribution (2 positive, 2 negative, 1 ground)
 - Assume standard ac cable, and dc current carrying ability is 67-percent higher than ac rating
 - Operational cost savings due to reduced energy use is not included
- ²⁵ Marquet, D., et al., New flexible powering architecture for integrated service operators, IEEE Intelec 2005
- ²⁶ Reliability numbers for rectifier from simulation of power train only. Will be lower with all components included. Assumes dc/dc converter between battery and bus – higher reliability will be achieved in systems without converter.