Intel Eco-Rack Version 1.5

Achieving 23-30% total power savings with a server rack designed for excellence in power efficiency[1]

White Paper

Intel® Eco-Rack

Executive summary

In 2007, Intel first demonstrated its Eco-Rack Proof of Concept (Eco-Rack 1.0). This was a real-world test of the power savings possible in rack-mounted servers using available technologies and making simple changes to existing components. Developed at Intel facilities in both Hillsboro, Oregon and DuPont, Washington, and inspired by discussions with the U.S. Environmental Protection Agency (EPA) and Lawrence Berkeley National Laboratory (LBNL), this demonstration showcased several industry advancements in energy-efficient technology for the data center. The result was a 16-18 percent total power savings based on workload or configuration.[2]

For 2008, Intel again assembled a test setup of 30 rack-mounted servers and tested some of the same, as well as some new, power management technologies and strategies for the data center. The result? The Intel Eco-Rack 1.5 achieved even better results: 23-30 percent total power savings, depending on the workload, as measured by running compute-intensive enterprise applications such as SPECjbb (a warehousing application simulating many transactions at the same time).[3]

In this paper, we describe what modifications Intel did to achieve these savings and make recommendations on which ones might be most effective. Some of the technologies and features we tested include:

- DC power delivery to servers
- Enhanced Intel SpeedStep® technology in the BIOS
- 2 GB memory modules
- Enhanced HALT state (C1E)
- Quad-Rank Fully Buffered DIMM (FB-DIMM)
- FB-DIMM Idle Power Enhancement Driver (FIPE)
- Memory refresh
- Closed Loop Thermal Throttling (CLTT)

Most surprising was that, depending on workload, CLTT alone accounted for 20.5 to 36.6 percent of total overall energy savings.

What makes all these savings so important is that data centers consume a measurable percentage of the world's energy—despite being such a relatively recent phenomenon in energy consumption. In an age increasingly concerned with carbon dioxide (CO2) emissions and global warming, achieving global goals for environmental protection, conservation and sustainability will require new technological advancements in many areas, including the data center. This innovative demonstration serves as an education tool to illustrate:

- The real-world possibilities inherent in using leading-edge technologies and employing various server-level configuration changes
- The energy savings and financial benefits these technologies and settings offer data centers
- The productivity gains energy-efficient data centers can deliver to businesses, governments, and other organizations

The future looks bright for power-saving developments such as the Eco-Rack. Intel and the industry are continuing to address the challenge of powering tomorrow's data centers and reducing their environmental footprint through a broad set of innovations that enable more work to be done with less power. As Francesco Serafini, VP Hewlett Packard EMEA, has said:

> *"ICT [information and communications technology] covers two percent of the global energy consumption. We can work to halve it to one percent, but the priority is to work to decrease the remaining 98 percent."[4]*

The message here is clear. As important as it is for ICT to cut its own consumption, it is even more important for ICT to continue providing valuable productivity gains which will reduce energy consumption in all sectors. These improvements will come in the form of everything from microprocessors in automobiles that improve gas mileage by monitoring every factor contributing to the efficacy of the fuel/air mixture, to chips in cell phones that may enable video conferencing and potentially reduce the number of plane flights taken for business.

Introduction

As economies around the world depend increasingly on digital information management, data centers housing the electronic equipment used for data processing, data storage, and communications networking have become essential to business, transportation, communications, government, and academic institutions. Many factors are driving the demand for more and larger data centers:

- Electronic transactions in financial services, such as online banking and electronic trading
- Internet communications and entertainment
- Electronic record-keeping for healthcare
- Global commerce and services
- Satellite navigation and electronic shipment tracking in transportation
- Emergency, health, and safety services
- Information security and national security
- High performance scientific computing
- Digital provision of government services, such as electronic filing of taxes and online postal tracking

The sheer number of people adopting and using Internet and related online services and applications is also driving data center growth. According to Jupiter Research's "Worldwide Online Population Forecast, 2006 to 2011," 1.1 billion people currently enjoy regular access to the Internet. This report goes on to predict that a compound annual growth rate of 6.6 percent will increase the number of people with regular online access to approximately 1.5 billion (or 22 percent of the Earth's population) in 2011.

Powering the data centers serving this growing number of online users is another matter. Data centers are a fast-growing phenomenon in the world of power consumption. According to the U.S. EPA, the nation's data centers and servers consumed about 1.5 percent of the U.S. total electricity consumption in 2006.[5]

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)."[6]

This same EPA report also suggests the problem is going to continue to grow, forcing data centers to compete for increasingly limited power resources with every other electrical user, from heavy industry to home users. The EPA predicts that "under current efficiency trends, national energy consumption by servers and data centers could nearly double again in another five years (i.e., by 2011) to more than 100 billion kWh, representing a $7.4 billion annual electricity cost" and the addition of 10 power plants to the nation's power grid.[7]

But competition for electricity is only one challenge data centers face. The cost of power is also troubling to many data center operators. According to Gartner, as power requirements continue to grow, energy costs will emerge as the second highest operating cost (after personnel) in 70 percent of worldwide data center facilities by 2009.[8] The average annual power costs for a 100,000 square-foot data center is nearly USD 6 million.[9] Some organizations are already cutting back on expansion because adding more servers would require a complete overhaul of the entire power and cooling infrastructure of their data center. This is in line with a prediction Gartner made back in 2006 that 50 percent of data centers would have insufficient power and cooling capacity by 2008.[10]

Microprocessors are an important part of the solution

What is surprising is that these problems are appearing despite some of the greatest energy efficiency gains per unit of processing power the world has ever seen. In March 2008, for instance, Intel introduced two energy-efficient 50-watt server processors that directly benefit companies with power-constrained, high-compute density environments. These Intel® Xeon® L5400 processors are as much as 25 percent faster[11] and have a 50 percent larger cache size than Intel's previous-generation, low-voltage Intel® Xeon® processors, while at the same time maintaining the low 50-watt thermal envelope. These processors use Intel's hafnium-infused high-k metal gate transistors, an innovation that has enabled Intel's quad-core 45nm low-voltage server chips to attain new heights in power-efficient performance.

Unfortunately, many processor efficiency advances like these are not being capitalized on yet because most data centers still run older equipment. A former director of Intel's IT operations admits that is true even at Intel, but that the company is moving faster now in taking advantage of its processor advances. He was quoted as saying that "a six-year-old server takes up valuable resources that could be better used, so we [Intel] have accelerated our refresh rate. Refreshing one data center gave us [Intel] three times the performance..."[12]

Saving energy in data centers requires more than energy-efficient processors

Replacing older servers with new ones using energy-efficient processors is just one part of the solution. A data center with all its systems for power conversion, distribution, backup storage, and cooling has many other options than just its computing power (processors) for saving energy. So, while much has been done to improve processor performance and efficiency, Intel and other industry leaders believe it is time to take a more holistic approach to planning, designing and laying out the data center to optimize power and cooling capacity. This includes finding ways to optimize all the variables, including building type, building systems, rack configuration, power conversion and distribution, monitoring and managing workloads, airflow dynamics, and many other energy-consumption factors.

Intel and the industry are addressing these variables through a broad set of innovations designed to enable data centers to do more work with less power. One of Intel's ongoing efforts is the Eco-Rack Proof of Concept. Eco-Rack Version 1.5 is meant to be a model of what can be achieved today using available technologies and making simple changes to existing components.

This paper will describe the current Eco-Rack's objectives, setup, and results.

The importance of rack-mounted servers

To understand why Intel chose to focus on server racks, it is important to understand their advantages in the modern data center. Today's high-density rack-mounted servers reduce space, cooling, networking, cabling and management costs—as well as total power consumption—compared to low-density configurations having far few servers per rack. However, some data center managers fear rack-mounted servers will drive up power density to levels exceeding the limits of their facilities.

While the power density of a fully-populated rack of 1U rack-mounted servers may be too high for most data centers,[13] many advantages of these racks—such as the lower power consumption per server—can be gained at lower densities. In fact, the primary total cost of ownership (TCO) benefit related to their use comes from their reduced power consumption, not their reduced space consumption. The use of more efficient cooling designs that concentrate just on the racks (rather than cooling the entire facility) can deliver additional savings. Recent research suggests that by installing rack-mounted servers at various densities, data centers could realize greater TCO benefits.[14]

Objectives for the Eco-Rack Proof of Concept 1.5

The Intel Eco-Rack Proof of Concept is a real-world example of the power savings and environmental footprint reduction possible in today's data centers through the use of rack-mounted servers employing today's power management technologies and making simple changes to existing components. Developed at Intel facilities in Hillsboro, Oregon and DuPont, Washington, the demonstration showcases industry advancements in the delivery of energy-efficient technology for the data center. It is designed to serve as an education tool to illustrate the real-world possibilities inherent in leading-edge technologies, the financial benefits their energy savings provide data centers, and the productivity gains that more energy-efficient data centers can deliver to businesses, governments, and other organizations.

With the Eco-Rack Proof of Concept Version 1.5, Intel shows how through relatively modest modifications to today's best-in-class server rack components, the power consumption of rack-mounted servers can be further reduced. Many of these modifications are power management features that are not used in the belief that they will affect performance or availability. In most cases, the effect of these features on performance and availability is negligible, while their potential for energy savings is considerable.

It should be noted that any reduction of power consumed in a data center has considerable ramifications in reducing heat and thus cooling costs. Every watt of power directly consumed by silicon in a typical data center requires another watt for power conversion (based on standard 75 percent efficient power supplies and 88 percent uninterruptible power supplies) and another watt for cooling the two watts of load (one watt, silicon; one watt, power conversion).

Test setups

To ensure a fair comparison, Intel set up two test racks—the Eco-Rack and a typical Modern AC Configuration—each holding 30 1U rack-mounted servers. The typical Modern AC Configuration provided the baseline for typical power consumption upon which the savings of each Eco-Rack modification was computed.

Each server featured two Intel® Xeon® processors X5355 (2.66 GHz). All servers were running Fedora Core 5* 64-bit with Kernel 2.6.15-1.2054_FC5 with IPMI enabled. The power factor for the uninterruptible power supplies (UPS) on the standard rack setup—referred to as the Modern Alternating Current (AC) Configuration in this paper—was 0.99. This 0.99 power factor was equivalent to the UPS performance of the direct current (DC)-based Eco-Rack setup, which, because it is a DC-based system, required no complex power phasing.

The goal was to compare a typical modern configuration with an Eco-Rack design providing equal performance, but requiring less power.

Below we further describe each configuration.

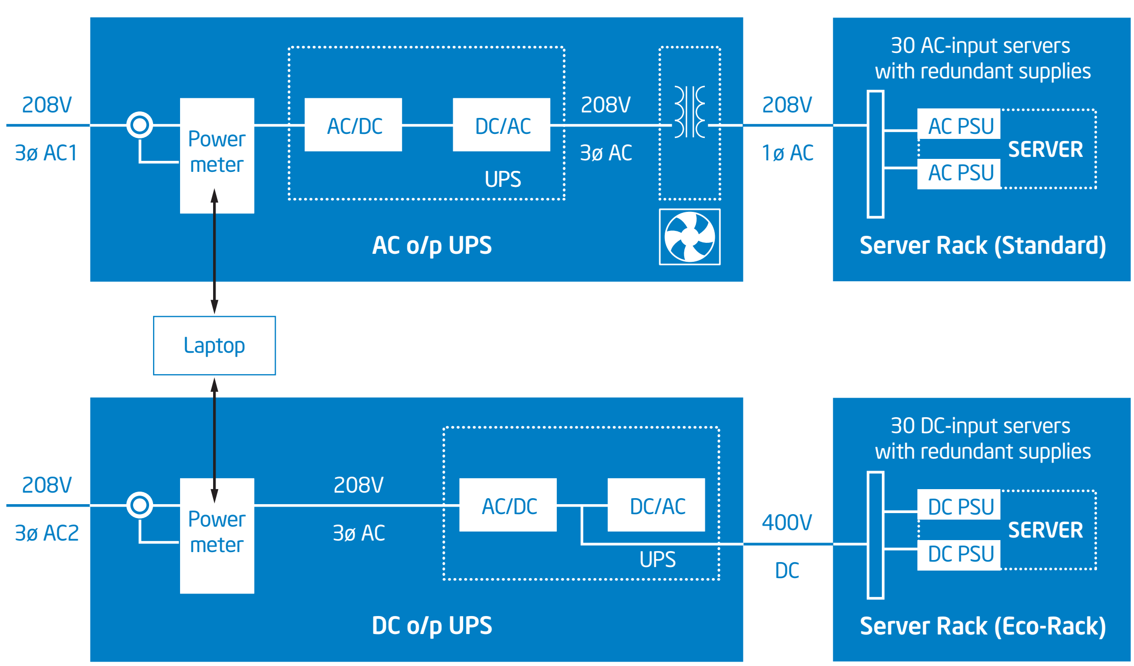

Figure 1. Test setup. Modern AC configuration vs. Eco-Rack setup. Note the fewer conversion stages in the Eco-Rack.

Modern AC configuration

This setup featured 30 1U rack-mounted servers equipped with Intel Xeon Processors X5355 (2.66 GHz, Stepping 11) and 8 GB of memory supplied by eight 1 GB FB-DIMMs. Supported by the Intel® 5000 chipset family and FB-DIMM technology, Quad-Core Intel Xeon processor 5300 series-based platforms deliver exceptional performance with better power efficiency: Up to 1.5 times the performance compared to leading Intel® Xeon® processor 5100 series in the same power envelope and up to 2.4 times the performance compared to the best published results of AMD Opteron* Model 2222.[15]

To provide the best possible setup for comparison with the Eco-Rack, we configured this modern rack with best-in-class components that are commercially available today. This includes using greater than 90 percent efficient power supply units (PSUs) for both the main and redundant PSUs for each system. To make the Modern AC Configuration the best baseline for our Proof of Concept, we turned off most of the selectable processor and chipset power management technologies. To provide the most accurate reading on power draw, we measured current coming directly from the wall rather than after it passed through a UPS.

In a typical data center, 480 volts is stepped down through a power distribution unit (PDU) to 208 volts AC to feed each server in the rack. To simulate this using 208 volts AC, we step across 208 AC through a PDU after it passes through the UPS. The PSUs, and later the voltage regulators, in each server then progressively convert and step down the current to the 12 volts and less DC required by the various server components. While the typical power conversion efficiency for the conventional architecture is just below 50 percent, a system using best-in-class components like this one can achieve up to 71 percent efficiency.

Eco-Rack Configuration. This setup started out exactly the same as the typical Modern AC Configuration. We then made a series of changes (described below) to evaluate the energy savings potential of each modification.

Eco-Rack modifications

In this section, we describe each of the modifications we made and the sequence in which they were done to try to improve energy efficiency. The effect of each modification was measured on four different workloads: idle, two active cores (25 percent total capacity), four active cores (50 percent), and eight active cores (100 percent).

1. Activating Enhanced Intel SpeedStep® technology

Enhanced Intel SpeedStep technology enables Intel® Server Boards to dynamically adjust processor voltage and core frequency, which can result in decreased average power consumption and decreased average heat production for a reduction in operating costs. By turning on Enhanced Intel SpeedStep technology through the BIOS, data center managers can reduce server power consumption and cooling costs with little effect on performance.[16]

As our first modification, Enhanced Intel SpeedStep technology proved very effective. Depending on workload, it accounted for 9.4 to 22.5 percent of total overall energy savings.

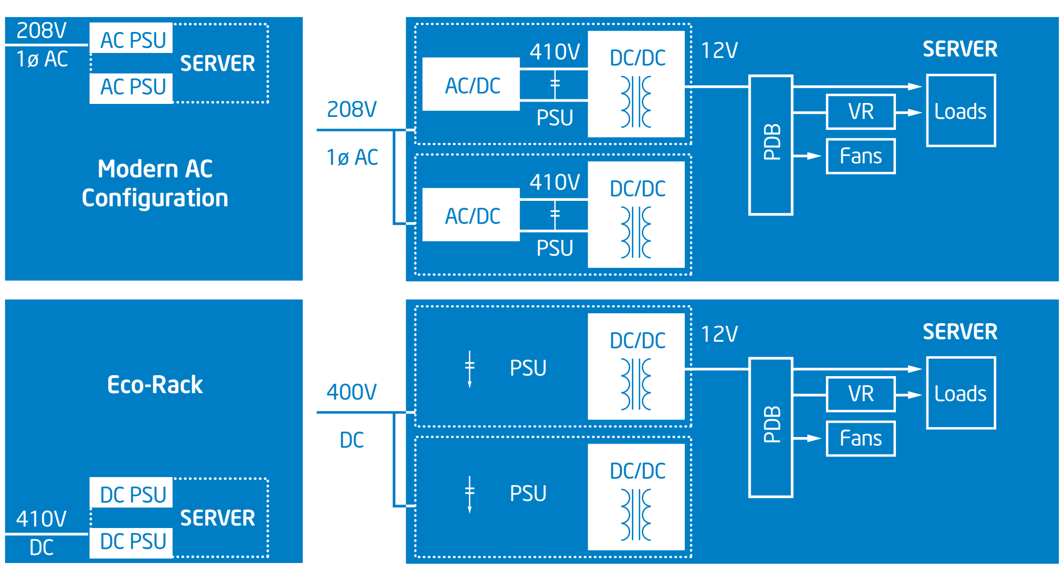

Figure 2. Server detail. These diagrams provide the power distribution details within each server. Note the fewer conversion stages in the Eco-Rack.

2. Power delivery

The second modification we tried was converting the Eco-Rack setup to DC power. We eliminated the initial power conversion step by using 400 volts DC power directly. This improves energy efficiency by reducing the power losses due to: 1) power conversion, and 2) the need to cool the excess heat that this conversion generates.

To use DC power, we deactivated the DC-back-to-AC power conversion in the UPS and sent the 400V DC directly to a DC power strip for the servers. Servers were equipped with DC PSUs to progressively step down the current to the 12 volts and less DC required by the various server components. Again, to provide the most accurate reading on power draw, we measured current coming directly from the wall rather than after it passed through a UPS.

Depending on workload, DC power delivery accounted for 14.1 to 33.8 percent of total overall energy savings.

3. Memory configuration

Our third modification involved memory configuration. Instead of the eight 1 GB FB-DIMMs used by the Modern AC Configuration, we inserted four 2 GB FB-DIMMs. Using fewer memory modules (four versus eight) reduces active power consumption, produces less heat, and requires less energy to be spent in cooling. It is important to note that for optimum performance, FB-DIMM memory modules should be equally divided between the memory channels (e.g., one memory module for each of four memory channels).

Depending on workload, the switch to 2 GB FB-DIMMs accounted for 11.2 to 35.3 percent of total overall energy savings.

(Note: Besides capacity, DIMMs can differ in wattage draw, memory speed, and the number of ranks. All of these can affect both energy consumption and performance. It is smart to compare carefully before making a purchase to ensure goals in energy savings and performance will be met.)

4. Enhanced HALT state (C1E)

C1E reduces processor speed and voltage when the processor is in idle mode, decreasing power consumption and heat generation. The C1E state is invoked when the operating system's idle process issues a HALT (HLT) command. Most operating systems do this repeatedly when not under a full load. The C1E halt state turns down the entire CPU's clock frequency (via multiplier control) and voltage, squeezing out extra energy savings.

Whether or not the processor actually gets the HLT command is determined by settings related to the chipset. Values are generally set in the BIOS. The setting/value can be enabled or disabled.

Some motherboards BIOS default to preventing the HLT command from being effective. Other motherboards do not even have an option in the BIOS setup to alter the defaults. In such cases, it may be possible to set values through third party software that manipulates the chipset registers.

(In most servers, enabling power management through the BIOS activates all native power management strategies, including C1E, Enhanced Intel SpeedStep technology, and others. Our setup allowed us to isolate each technology to evaluate their effect.)

Depending on workload, setting the C1E state to "enabled" versus "disabled" accounted for between 4.5 to 15 percent of total overall energy savings.

5. Quad-Rank 2 GB FB-DIMMs

For our fifth modification, we tried exchanging the 2 GB FB-DIMMs used in our third modification (described above) with low-power quad-rank 2 GB FB-DIMMs. Most 2 GB FB-DIMMs are standard two-rank registered memory. They are high density FB-DIMMs created using high-end chips arrayed across the memory board. Quad-rank FB-DIMMs are twice as tall, with the top set of chips flipped over 180 degrees to reduce latency between top and bottom. This enables production of 2 GB FB-DIMMs composed of four 512 MB memory chips. These new quad-rank FB-DIMMs include a low-power advanced memory buffer. Our intent was to reduce the wattage footprint of the Eco-Rack's memory modules while also using a lower cost memory product. As before, it is important to note that for optimum performance, FB-DIMM memory modules should be equally divided between the memory channels.

This modification did not achieve significant results. In fact, under most workloads it actually was less efficient, accounting for -3.8 to 4.6 percent of total overall energy savings. Only at idle workloads did it have a positive effect. However, if performance is a priority, quad-rank is the way to go.

6. Fully Buffered DIMM Idle Power Enhancement Driver (FIPE)

The sixth modification in our progression evaluated FIPE, a software driver available only for Linux.* This software runs continuously querying sequentially each of the four cores in each socket to check how busy they are. If they are not busy, it throttles the memory subsystem down to a lower frequency to save power. We activated this feature by compiling and turning the driver on (insmod method).

This modification had its greatest effect at idle and 25 percent workload (3.6 percent and 5.7 percent of total energy savings, respectively). It had zero effect at a 50 percent workload, but then managed 2.7 percent of total energy savings at a full workload.

If we average this, anticipating that the workload will vary throughout the day (which is the case with most data center servers), we could anticipate total overall energy savings of about 3 percent for this modification.

This feature will be in future Linux kernels, and is included natively in hardware such as dual-processor platforms with Intel® 5400A or 5400B chipsets.

7. Memory refresh

Memory refresh is the process of periodically reading information from an area of computer memory and immediately rewriting the read information to the same area with no modifications. Each memory refresh cycle refreshes a succeeding area of memory. Memory refresh is most often associated with modern dynamic random access memory (DRAM). Double memory refresh is the default. For this test, we changed the setting to single memory refresh (forcing it on through the operating system—though in some systems it can be set in the BIOS). Instead of memory being accessed two times for every refresh cycle, it was now only accessed once. The assumption here was that by reducing the number of times memory is accessed for a refresh, we would save power.

This assumption turned out to be wrong. In fact, at 25 percent workload, single memory refresh resulted in a 15.6 percent loss in total overall energy savings. At other workloads, its effect was less dramatic, and at 50 percent workload there was actually a slight improvement (0.7 percent) in total overall energy savings.

8. Closed Loop Thermal Throttling (CLTT)

CLTT is a feature of the Intel® chipset to prevent FB-DIMM memory from overheating. This is a temperature-based throttling feature used with variable speed fans. If the temperature of the installed FB-DIMMs approaches their thermal limit, the system BIOS will initiate memory throttling. This limits bandwidth to the DIMMs, therefore capping the power consumption and preventing the DIMMs from overheating. By default, the BIOS will configure the system to support CLTT if it detects that there are functional Advanced Memory Buffer (AMB) thermal sensors present on all installed DIMMs (though sometimes the sensors are on the motherboard, or both). In CLTT mode, the system fans run slower to meet the acoustic limits for the given platform, but will also allow the fans to ramp up as needed to maintain the parts within temperature specifications under high stress levels. In a similar way to Open Loop Thermal Throttling (OLTT), the system BIOS utilizes a Memory Reference Code (MRC) throttling algorithm to maximize memory bandwidth for a given configuration. The MRC code relies on Serial Presence Detect (SPD) data read from the installed DIMMs.

Closed Loop Throttling is autonomous and internal to the MCH hardware and the throttling algorithm always throttles the minimum amount required to prevent the part from overheating. CLTT is the Intel preferred platform control mechanism as it provides the best memory bandwidth performance while providing the lowest system fan acoustics.

CLTT turned out to be very effective in reducing power consumption. Depending on workload, CLTT accounted for 20.5 to 36.6 percent of total overall energy savings.

The tests

Tests were done under the guidance of the Lawrence Berkeley National Laboratory in California. The test configurations ran JILL5, a version of the industry-standard SPECjbb benchmark that tests Java Virtual Machine* performance by running a Java-based order fulfillment warehouse application. SPECjbb is generally regarded by the industry as being a realistic JVM performance benchmark with a real-world typical instruction mix. JILL5/SpecJbb does not measure I/O performance (e.g., disk, network) since all its transaction requests are synthesized internally.

One of the configuration parameters is the number of warehouses to run in parallel. For power-measurement purposes, we did separate runs at two, four and eight warehouses to approximate 200, 400, and 800 percent of normalized single-CPU utilization.

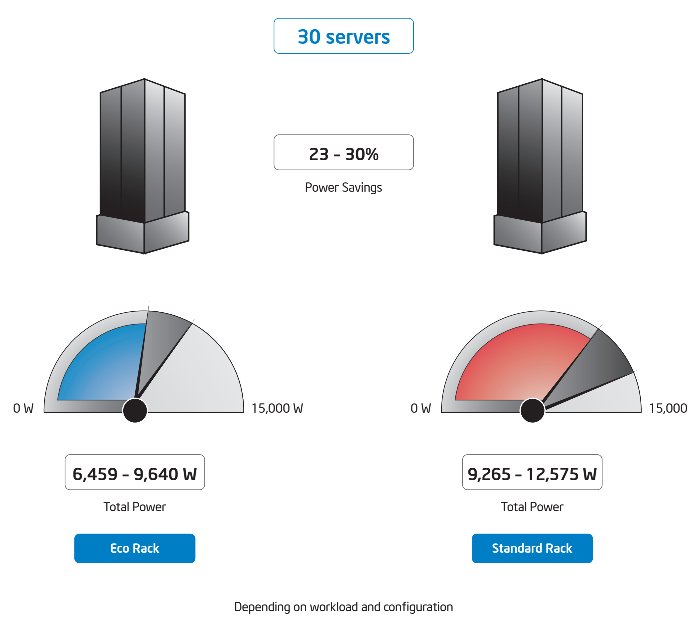

Figure 3. Power comparison.

Since JILL5/SPECjbb is intended to be a benchmark, it is normally run for a fixed amount of wall clock time, after which it reports a result in billions of operations per second (bops). Since we are using it to simulate a continuous system load instead of as a benchmark, we set its run-time to 9...999 minutes instead of the normal benchmark runtime. We never wait for it to finish.

Note that as the number of warehouses-in-parallel approaches the number of logical processors, the CPU utilization drops a bit below the theoretically attainable amount of 100 percent of X number of warehouses. This is probably due to the warehouse threads contending with one another for some single resource such as a lock on the results-counter, overall memory contention, other CPU resources, etc.

Test results

Running JILL5/SPECjbb, the Eco-Rack demonstrated measured total power savings of 23 to 30 percent, depending on workload or configuration, over a standard data center AC server rack configuration.

Projecting even the lowest power savings (23 percent) to a data center with 1,000 servers and a power cost of USD 0.062 per kilowatt hour, the modest changes represented by the Eco-Rack configuration could annually save a company or organization tens to hundreds of thousands of dollars depending on the size of the data center, the number of racks, and the number of servers per rack.

In a power-constrained environment, these same power savings could also be used to make better use of available power and add more compute cycles. (See hypothetical example below.)

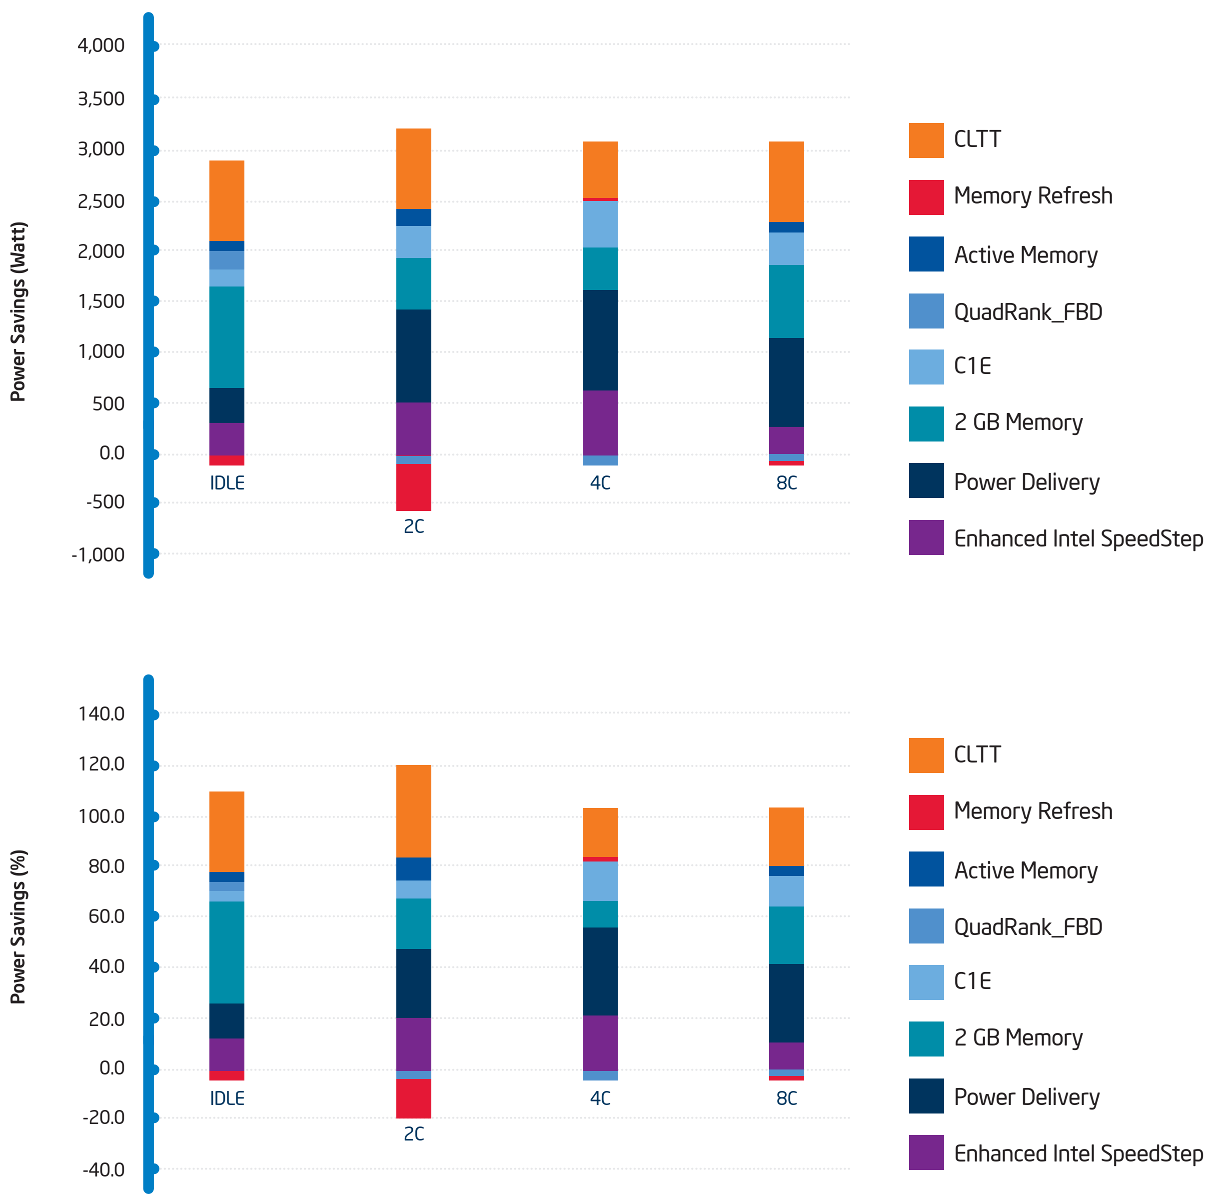

Figure 4. Power stack-up: Building from the baseline down.

A hypothetical example of how data center power constraints affect operations

To see the advantages of something like the Eco-Rack and other data center energy efficiency enhancements, it is helpful to look at a hypothetical example of how the power constraints of a data center might actually hurt a company's bottom line.

Consider a financial firm that is faced with a data center that has outgrown its power resources. Like many financial companies, this firm has to borrow a certain amount of money at the end of every day to hedge its next day's loans and other money it needs to have available. If the firm does not have enough power in their data center to run the compute cycles it needs to accurately compute how much it needs to borrow, the firm has to increase this hedge, borrowing more than necessary to be "safe." Many financial companies do this, borrowing and paying for more money than they need to make up for their inability to accurately and quickly calculate what they need for the next day. This directly affects a firm's profits. By updating their data centers with modifications such as those demonstrated by the Eco-Rack Version 1.5, financial firms could make better use of their available power and potentially gain the compute cycles they need to much more accurately compute their daily hedge. For many, the interest saved on borrowed hedge money would far outweigh the cost of the power to run these compute cycles on the more efficient servers.

Conclusions

The Eco-Rack Proof of Concept demonstrates how server-level changes—such as converting to DC power delivery to servers, activating Enhanced Intel SpeedStep technology in the BIOS, using more efficient 2 GB memory modules, and activating CLTT—can deliver significant savings at the rack-level for today's data centers. Incorporating these changes resulted in a measured total power savings of 23 to 30 percent, depending on workload, when running compute-intensive enterprise applications such as SPECjbb, a warehousing application simulating many transactions at the same time.

Many of these "small" changes are things data center managers can do today. In fact, if we could make just one recommendation from this study, it would be to activate CLTT. Depending on workload, CLTT alone accounted for 20.5 to 36.6 percent of total overall energy savings.

Utilizing power optimization features and settings can go a long way to saving energy and reducing TCO. So can specifying best-in-class, energy-efficient components in everything from server processors to PSUs. Other changes, such as DC power components, will take more time for the industry to adopt. One can also save from 30 to 50 watts per server, depending on workload, by removing the redundant power supply.

The future looks bright. Intel and the industry are continuing to address the challenge of powering tomorrow's data centers and reducing their environmental footprint through a broad set of innovations that enable more work to be done with less power. We summarize many of these below.

Promising power-reduction technologies and strategies for data centers

400V DC Distribution

Having each server in the rack accept 400V DC reduces the number of power conversion stages compared to the Modern AC Configuration. The AC-to-DC conversion in the UPS, the AC-to-AC conversion in the PDU transformer, and the AC-to-DC conversion in the PSU are eliminated. Fewer power conversion steps reduce energy lost to conversions and result in fewer power conversion components to cool. In addition, since fewer components are required for power conversion, system reliability is expected to increase.

Distribution at 400V DC also allows for a more efficient connection of renewable energy sources to a building since most renewable energy sources (e.g., fuel cells and photovoltaic panels) generate a DC voltage. The use of DC power also eliminates the need for balancing the load across the three phases as required in AC distribution. This saves time and effort and eliminates the need for equipment monitoring the load imbalance.

Overall, the advantages of direct high voltage DC systems (400VDC) include:

- Removes redundant conversion in end-to-end data center power architecture
- Seven- to 10-percent energy-efficiency gain
- 10 times higher availability[17]
- Commonly available UPS, battery and PSU technology (just less parts)
- Lower thermal footprint
- No complex power phasing/balancing
- Easy-to-parallel power sources

High Performance, Energy-Efficient Quad-Core Processors Based on Intel® Core™ Microarchitecture

For every watt saved in computation, two additional watts are saved—one watt in power conversion and one watt in cooling (the result of no longer having to cool two watts in computation and power conversion). Consequently, big power savings can be gained from small percentages in processor power savings. Information on the latest gains in performance and energy efficiency by Intel Xeon processors can be found at: www.intel.com/products/server/processors/index.htm.

Processor Voltage Optimization

Intel Core microarchitecture's Intel® Intelligent Power Capability throttles down processor performance in response to lower demand. This capability enables excellent energy efficiency in a wide variety of loads. Data centers using the latest Intel processors will profit from this technology.

Optimizing BIOS and System Firmware Settings

Through Enhanced Intel SpeedStep technology, Intel Server Boards can dynamically ramp processor speed and voltages to minimize power consumption and reduce operating costs. By turning on Enhanced Intel SpeedStep technology through the BIOS, data center managers can reduce server power consumption and cooling costs by up to 25% with little effect on performance.[18] Updating system firmware settings to appropriately reflect system components, such as processor number and type and fan speeds, can often reduce the speed at which fans run, providing additional power savings.

Greater Than 90 Percent Efficient PSUs

Typical server PSUs (one element in a data center's power conversion chain) are only 50 percent efficient, yet units can be purchased that are 90 percent efficient or more and pay for their additional cost in less than a year. Depending on the number of servers in a data center, 90 percent efficient PSUs can have a significant effect on power savings.

Improved Rack Design

Cooled doors, chimneys, raised floors, and other data center design improvements increase air flow effectiveness and reduce cooling costs. Much information is already available from Intel and other sources on ways to design data centers so cooling is directed at the components putting out the heat and not at the facility at large.[19]

Server Consolidation

In many data centers, more servers are running than necessary for the work being done. A study by Hewlett Packard Lab of six corporate data centers studied 1,000 servers and found that most were using only 10 to 25 percent of their capacity.[20] To combat this trend, a major Intel focus is promoting better server utilization through server consolidation. Consolidation through virtualization reduces the number of physical servers required, thereby helping to minimize floor space, cooling, and capital costs. Many data center operators are responding to the call for consolidation by installing virtual machine software on powerful Intel®-based servers equipped with Intel® Virtualization Technology (Intel® VT). This enables a single server to run multiple operating systems and applications.

An important benefit of server consolidation is greater energy efficiency. Having fewer servers do more work reduces both the number of servers drawing power and the number of servers that need cooling.

Intel IT did a benchmarking study comparing two- and four-socket servers with Intel® quad-core processors in a virtualized environment (see: http://communities.intel.com/docs/DOC-1512). The study found that in memory-intensive deployment scenarios, a four-socket server based on the Intel® Xeon® processor 7300 series can support approximately 15 percent more virtual machines for the same TCO than a deployment based on two-socket servers. For scalability-focused scenarios, the four-socket server offered nearly twice the scalability. The Quad-Core Intel Xeon processor 7300 series provides up to 1.9x scalable performance improvement over two-socket Intel Xeon processor 5300 series-based servers.[21] Only in situations where server demand was limited by business policy or other factors was the two-socket server (based on the Quad-Core Intel Xeon processor 5300 series) more cost-effective.

For those situations such as license restrictions or more modest virtualization workloads, the Intel® Xeon® processor 5300 and 5400 series are excellent choices. This processor series delivers four energy-efficient processor cores and Intel VT to deliver the most headroom for virtualization on a two-socket server, all while providing record-setting performance. (See performance site for full details: www.intel.com/performance/server/xeon/vt_vcon1.htm.)

Some of the acronyms used in this paper

AC	Alternating Current
AMB	Advanced Memory Buffer
CLTT	Closed Loop Thermal Throttling
C1E	Enhanced HALT State
DC	Direct Current
DRAM	Dynamic Random Access Memory
EPA	U.S. Environmental Protection Agency
FB	Fully Buffered
FIPE	Fully Buffered DIMM Idle Power Enhancement Driver
HLT	HALT
ICT	Information and Communications Technology
Intel® VT	Intel® Virtualization Technology
LBNL	Lawrence Berkeley National Laboratory
OLTT	Open Loop Thermal Throttling
PSU	Power Supply Unit
SPD	Serial Presence Detect
TCO	Total Cost of Ownership
UPS	Uninterruptible Power Supply

Learn more

Visit Intel's Technology & Research site at www.intel.com/technology

Find out about Intel's Eco-Technology Innovation at www.intel.com/technology/index3.htm

Download the Intel white paper "The State of Data Center Cooling" at
http://download.intel.com/technology/EEP/data-center-efficiency/state-of-data-center-cooling.pdf

Download the Intel white paper "Turning Challenges into Opportunities in the Data Center" at
www.intel.com/technology/eep/data-center-efficiency/dcwhpaper.pdf

1. Based on workload or configuration.
2. "Intel® Eco-Rack: Achieving 16-18% total power savings with a server rack designed for excellence in power efficiency," Intel white paper, 2007.
3. SPECjbb2005* is slightly different than SPECpower_ssj2008* (a simulation we intend to use for future benchmarking).
4. "EU plans mandatory energy-efficiency standards for ICT," EurActiv.com, February 22, 2008.
5. *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.
6. Ibid.
7. Ibid.
8. *Gartner Says 50 Percent of Data Centers Will Have Insufficient Power and Cooling Capacity by 2008,* Gartner Inc. press release, November 29, 2006.
9. Mark, Roy, *House Green Lights EPA Data Centers Study,* Internetnews.com (July 13, 2006)
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. 25% claim based on performance comparison on SPECpower*_ssj2008 benchmark. Comparison details: Baseline: SuperMicro X7DBE+ Server platform with 2 x Intel® Xeon® processor L5335 2.0 GHz, 8 M L2 Cache, 1333 MHz FSB, Intel® 5000P chipset, 4x2 GB FBDimm 667 MHz memory, Microsoft Windows Server 2003 Enterprise x64 Edition SP1,* BEA JRockit JVM.* New platform measurement: Same configuration as the baseline except with 2x Intel® Xeon® processor L5420 2.50 GHz, 12 M L2 Cache. Intel internal measurement. March 2008.
12. *Intel Pushes Denser Data Centres,* The Register (February 24, 2007).
13. 1U is a unit of measure used to describe the height of a server or other device mounted on a rack. 1U is 1.75 inches high.
14. Source: *Cooling Strategies for Ultra-High Density Racks and Blade Servers,* APC White Paper #46, by Neil Rasmussen: www.apcmedia.com/salestools/SADE-5TNRK6_R4_EN.pdf
15. 1.5 times or 50% performance claim based on performance with measured SPECint*_rate_base2000 benchmark results comparing Intel® Xeon® Processor 5345 and Intel® Xeon® Processor 5160. Published/measured results as of Sept 21, 2006. 2.4 times or 140% over Opteron performance claim based on published results on SPECjbb*2005 benchmark on Intel® Xeon® processor 5365 and AMD Opteron* Model 2222SE. Results are best published as of August 7, 2007. For more information about server performance, visit www.intel.com/performance/server.
16. Source: Intel White Paper: *With Intel® Server Technologies Built In, You Have Flexibility and Savings Built In,* 2006.
17. Marquet, D., et al., "New flexible powering architecture for integrated service operators", IEEE Intelec Conference, 2005. Also, T. Yamashita, et al., "270 V DC – A Highly Efficient and Reliable Power Supply System for Both Telecom and Datacom Systems," IEEE INTELEC Conference, 1999.
18. Source: Intel White Paper: *With Intel® Server Technologies Built In, You Have Flexibility and Savings Built In,* 2006.
19. See Intel White Paper: D.G. Costello, P. F. Grimm, M. Loeffler, M.K. Patterson, "Data Center TCO: A Comparison of High-density and Low-density Spaces" (*Thermes,* 2007)
20. A. Andrzejak, M. Arlitt and J. Rolia, "Bounding the resource Savings of utility Computing Models," working paper HPL-2002-339, Hewlett-Packard Laboratories, Palo Alto, California, November 27, 2002. as cited in Nicholas G. Carr, "The End of Corporate Computing," MIT Sloan Management Review, Volume 46, Number 3, Spring 2005, pp. 67-73.
21. Performance measured using the TPC-C* benchmark. Intel internal, early platform measurements (June 2007) comparing system configurations of four Quad-Core Intel Xeon processor 2.66GHz or 2.93GHz, with 2x4M L2 Cache, 1066 MHz system bus, Clarksboro Chipset with two Intel® Xeon® processor X5355, 2.66 GHz with 2x4M L2 Cache, 1333 MHz system bus, Blackford Chipset.

Printed in USA 0808/VP/HBD/PDF

Please Recycle