



White Paper
Energy Efficiency in
the Data Center

Turning Challenges into Opportunities in the Data Center

Contents

Introduction.....	3
In Search of a True Yardstick for Data Center Performance.....	3
Focusing on the Entire Data Center	3
Driving Energy Efficiency at the Processor Level.....	4
Introducing 45nm High-k next-generation Intel® Core™ microarchitecture.....	4
Driving Energy Efficiency at the Platform Level.....	5
Server consolidation	5
More efficient power conversion	5
Defining the “energy-efficient server”	5
Power Management Bus (PMBus)	6
Driving Energy Efficiency at the Data Center Level	6
Researching the most optimum power delivery system for data centers.....	6
Improving data center design.....	7
Promoting incentives for data center energy efficiency.....	7
Research at Intel’s own data centers.....	7
What Intel’s Efforts Mean for the Data Center of the Future.....	7
Learn More	8

Introduction

Data centers are an integral part of the engine driving today's economy, but they're becoming increasingly more expensive to power and cool. Currently, data centers consume a little over one percent of the power in the United States annually. Some large data centers draw as much energy as a small town. These figures continue to rise as compute needs grow, density increases, and power and cooling demands climb. As a result, many data centers are reaching their full capacity sooner than expected.

The picture isn't as dark as it seems. Intel sees many opportunities to improve energy efficiency in nearly every aspect of data center operation that consumes power. In fact, in addition to our industry-leading two-year cadence for developing a new silicon process technology and microarchitecture for delivering processors that provide greater "performance per watt," Intel is taking a leading role in working on solutions for reducing and managing power throughout the data center. Working collaboratively with industry and government organizations, Intel is providing products, technologies and ecosystem leadership to help accelerate the industry toward the goal of optimizing the ratio of compute output to total power consumed by the data center.

In Search of a True Yardstick for Data Center Performance

The first requirement for improving energy efficiency in the data center is the development of reliable metrics for comparison of energy-efficiency performance. This has been problematic as previous metrics have concentrated on performance per watt at just single points along the loadline — such as all the data center's servers at idle or at 100% utilization. Neither of these points are an accurate measure of real-life usage since the amount of work each server does can vary greatly depending on its function and usage. A more reliable metric is to measure the entire loadline of all the servers for a day, week and month, determining the total compute output as a ratio of total power consumed in the entire data center.

Intel is working with industry-standards bodies such as the Standard Performance Evaluation Corporation (SPEC) and The Green Grid to determine performance and energy benchmarks that are comprehensive and unbiased. Reliable industry-accepted standards would give server manufacturers a better guideline for manufacturing energy-efficient servers and a better way of demonstrating

their effectiveness to customers. Such standards could also promote more competition in energy efficiency between vendors as they seek to post better scores on the benchmark tests.

Focusing on the Entire Data Center

Intel was recently recognized for its efforts in reducing data center power consumption through its multi-core server processors and Intel® Core™ microarchitecture, a combination that has greatly improved performance per watt and efficiency per square foot. While watts per square foot have risen, Intel has enabled performance per watt to rise even faster, enabling more work to be done per square foot. Today, for example, through Intel® processor improvements, 53 blade servers drawing 21 kW provide the same 3.7 Tflops of performance that in 2002 would have required Intel® Architecture-based HPC clusters using 512 servers arranged in 25 racks using 128 kW of power. That means just 17% of the power required in 2002 now provides the equivalent performance.

Addressing power consumption at the compute level pays rich dividends. Every compute watt reduced has a multiplier effect of three. 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.

To add to these power savings while continuing to focus on performance per watt efficiency at the processor level, Intel is helping the industry address the other two watts — the ones consumed in power conversion, cooling, and lighting and everywhere else the data center uses power. To do this, Intel is concentrating on a number of vectors:

- **Power management** — controlling the platform power based on actual workload and minimizing the power consumption not associated with computing.
- **Power reduction** — focusing on fine-grained power management capabilities to improve efficiency through lower power states and workload-tunable power states (being able to turn off components and systems not performing work and then bring them back when needed with as little latency as possible).

- **Power delivery** — looking for ways to minimize wasted power during conversion and transmission.
- **Regulatory influences** — promoting energy efficiency through industry-accepted metrics and incentives.
- **Enabling the ecosystem** — providing leadership and support for energy-efficiency efforts throughout the data center.

Driving Energy Efficiency at the Processor Level

Moore's Law is alive and well in the 21st century, continuing to enable Intel to deliver more performance in the same power envelope year after year. Today, a commodity server drawing roughly 400 watts delivers the same performance (roughly 100 Gflops) as the first supercomputer on the HPC Top 500 list in 1993.

In 2007, Intel continues its industry-leading two-year cadence for microarchitecture improvements and process technology. The start of the year saw Intel's announcement of its first 45nm processor featuring a new transistor breakthrough. Through a new material combination of High-k gate dielectrics and metal gates, Intel will be able to continue delivering record-breaking server processor speeds while reducing the amount of electrical leakage from transistors that can hamper chip and server design, size, power consumption, noise and costs.

Intel reaches energy-efficient milestone with 50-watt, high-performing quad-core server processors

In March 2007, Intel introduced two energy-efficient 50-watt server processors, the Quad-Core Intel® Xeon® processors L5320^A and L5310^A, which deliver a 35- to nearly 60-percent decrease in power from Intel's existing 80- and 120-watt quad-core server products. Requiring just 12.5 watts of power for each core, these processors provide similar performance yet set a new standard in energy efficiency. These new processors represent a nearly ten-fold improvement in power consumption per core in just 1½ years.¹

Intel's new 45nm silicon process technology will deliver greater levels of performance, energy efficiency and economies of scale.

- Double the transistor density
- Greater than 20% improvement in transistor switching speed or greater than 5x reduction in leakage power
- Greater than 30% reduction in switching power

Introducing 45nm High-k next-generation Intel® Core™ microarchitecture

In the second half of 2007, Intel is introducing the first family of 45nm processors. Codenamed Penryn, this family includes new Dual- and Quad-Core Intel® Xeon® processors. The Penryn family's microarchitecture optimizations will increase the overall performance and energy efficiency of the already leading Intel Core microarchitecture, delivering yet more instructions executions per clock cycle and speeding up the performance of a variety of applications.

Key advances include:

- Microarchitecture optimizations that increase overall performance and energy efficiency of the already leading Intel Core microarchitecture to deliver more instruction executions per clock cycle.
- Higher overall clock frequencies within existing power and thermal envelopes to further increase performance.
- Fast division of numbers for scientific computational workloads through the inclusion of a new, faster divide technique called Radix 16.
- Larger and faster caches to improve response time in accessing frequently used data.
- Intel® Streaming SIMD Extensions 4 (SSE4) instructions, the largest unique instruction set addition since the original SSE Instruction Set Architecture (ISA).
- Unique super shuffle engine that implements a full-width, single-pass shuffle unit that is 128-bits wide to perform a full-width shuffle in a single cycle. This significantly improves performance for certain content creation, imaging, video and higher performance computing.

Intel processors also offer additional features that help reduce power consumption in a wide range of hardware and software environments. A good example is Demand Based Switching (DBS). This feature dynamically tailors power consumption to workloads, ratcheting down processor power states whenever peak performance is not required. This can substantially reduce average power consumption for servers operating at typical data center utilization rates. Intel processors also feature dynamic power management. Power delivery to key processor subsystems is dynamically managed during runtime to continuously optimize performance versus power efficiency as workloads vary.

Driving Energy Efficiency at the Platform Level

At the platform level, Intel is leading the way towards greater energy efficiency through improving server utilization, promoting more energy-efficient platform components, and spurring energy-efficiency efforts throughout the server ecosystem.

Server consolidation

A major Intel focus is promoting better server utilization through server consolidation. Server utilization rates are often as low as 10 to 15 percent in data centers. 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.² Many data center operators today are installing virtual machine software on powerful Intel®-based servers equipped with Intel® Virtualization Technology¹ to run multiple operating systems and applications on a single server. Consolidation through virtualization reduces the number of physical servers required, thereby helping to minimize floor space, cooling, and capital costs. Increasing server utilization of servers through consolidation also improves energy efficiency. Intel's 45nm Penryn server processors will improve virtualized server performances through enhanced Intel Virtualization Technology. This innovation speeds up virtual machine transition (entry/exit) times by an average of 25 to 75 percent.

More efficient power conversion

One potential area for improving energy efficiency that Intel finds particularly promising is power conversion. For instance, the typical server power supply unit (one element in a data center's power conversion chain) is only 75-percent efficient, yet units can be purchased that are 90-percent efficient and pay for their additional cost in less than a year. Adoption of these efficient power supplies

has been slow since server purchases are typically done by IT managers as opposed to data center facility managers. IT managers tend to focus on getting the best purchase price, whereas a facility manager would look at the long-term power cost as well. As a result, small components that can have a big effect in the long run end up getting overlooked.

Defining the "energy-efficient server"

To encourage manufacturers to start using best-in-class energy-efficient power supply units and other components, as well as speed the adoption of energy-efficient servers in the data center, Intel is helping to define what exactly an energy-efficient server is. This is very similar to determining metrics for the data center as a whole, only at the server level. The metric needs to be based on agreed-upon workload applications, datasets and activity patterns consistent with real-life usage.

The critical importance of software

Multi-core processors deliver immediate performance and energy-efficiency benefits for multi-threaded applications, which include all major operating systems and database software, as well as many technical, financial and high-performance computing (HPC) applications. They can also improve total throughput for multiple applications running simultaneously in virtualized server environments.

Intel is working extensively with the industry to drive a broad migration to multi-threaded software (and to 64-bit software where appropriate). This includes providing:

- World-class software development tools, including Intel® compilers optimized for multi-threaded throughput and compatible with popular development environments.
- Optimized platform software that makes basic functions quicker and less energy-intensive.
- Deep collaboration with leading software companies to further improve performance and energy efficiency on the latest multi-core Intel processor-based platforms.
- Extensive training and support through the Intel® Software Network and Intel® Software College.

In general an energy-efficient server would:

- Reduce energy usage for non-compute functions through:
 - Energy-efficient power delivery (using energy-efficient voltage regulators, power supply units, and other components).
 - Energy-efficiency processors incorporating advanced power management capabilities and manufactured with the latest silicon process technology using High-k gate dielectric materials to minimize leakage.
- Provide fine-grained component power management (such as memory power management and Demand-Based Switching) along the load line to:
 - Minimize power consumption at system idle.
 - Provide monitoring and policy-based control of platform power consumption.
 - Allow peak server power consumption to be controlled externally during periods of high demand.

Regardless of how many energy-efficient features are used in a server, a server isn't recognizable to a purchaser as energy efficient without proper validation and information on the actual power draw and cooling load. To accomplish this, server manufacturers must provide an American Society of Heating, Refrigerating & Air Conditioning Engineers (ASHRAE) Thermal Report or equivalent. This report would give the data center owner and design engineer minimal, average, and peak power dissipation as well as airflows required by the server.

The U.S. Senate and House have approved bills to "give high priority to energy efficiency as a factor in determining the best value and performance of computer servers." The law has instructed the Environmental Protection Agency (EPA) to conduct a study analyzing the state of the art of data centers and servers in the U.S., including potential cost savings from the use of energy-efficient products. Intel is working with the industry to provide the best possible information to the EPA as it finalizes its study.

Power Management Bus (PMBus)

Intel is a member of the PMBus organization, the group responsible for this new open-standard digital power management protocol. The organization was founded by a coalition of power supply and semiconductor manufacturers who recognized that the lack of a suitable standard was inhibiting adoption of an all-digital power management solution.

The PMBus protocol is rapidly gaining wide industry acceptance. It facilitates communication with a power converter or other device by fully defining the transport and physical interface, and the command language, needed to accomplish this. This enables programming, control and real-time monitoring of power conversion devices. Having access to system power consumption data through a management controller enables OEMs and ISVs to develop innovative power management capabilities that solve real-world data center problems.

Driving Energy Efficiency at the Data Center Level

Intel is working within The Green Grid, a nonprofit consortium of information technology companies and professionals seeking to lower the overall consumption of power in data centers around the globe. The organization is chartered to develop meaningful, platform-neutral standards, measurement methods, processes and new technologies to improve energy-efficient performance of global data centers.

The Green Grid fills an important gap for the industry. While many standards bodies are looking at different aspects of data center efficiency, The Green Grid is the first organization that is uniquely focused on data center efficiency with an approach starting with the data center facility and working down to each individual area of data center efficiency.

Researching the most optimum power delivery system for data centers

In addition to its efforts in The Green Grid, Intel is also driving internal research. To find ways to improve efficiency throughout the power conversion chain, Intel has evaluated different power delivery architectures. A particular focus was finding the most efficient power architecture for operating at the wide range of loads experienced by the typical IT data center. Intel studied three different power delivery architectures – conventional alternating current (AC), rack-level direct current (DC), and facility-level DC distribution. Results showed the greatest efficiency is achieved through facility-level 380V DC distribution. Intel calculated that an efficiency of approximately 75 percent may be achieved with facility-level 380V DC distribution using best-in-class components. Compared to the typical data center with 50-percent power efficiency and drawing 10 MW for the computing infrastructure load, the 380V DC data center would draw only 6.67 MW. Assuming a cost of \$0.10 per kilowatt (kW) hour, the 380V DC architecture

could save approximately \$2.8 million in energy costs per year. Intel, along with several other industry partners, contributed to a small-scale demonstration of 380V DC facility-level distribution coordinated by Lawrence Berkeley National Lab.

Improving data center design

Another Intel® research project focused on the construction and design of data centers. Research showed that high-density data centers will provide the data center owner with a reduced cost of ownership when compared with that of a low-density data center. High-density data centers require specific design considerations, most notably a path for the higher volume of air. This could preclude some legacy data centers from moving to high density without local enhanced cooling, but new data centers and those with sufficient height for a retrofit can benefit from increasing densities. There are risks and ergonomic negatives to a high-density configuration, but these can be overcome by proper design and recognition that modern data centers do not require continuous staffing. If the data center can be designed or retrofitted to support the infrastructure for high-density computing, the owner will be able to have a smaller data center with the same computing performance at a lower total cost of ownership (TCO).

Promoting incentives for data center energy efficiency

Intel is also helping motivate energy efficiency in data centers through customer incentives. For example, Intel helped Northern California utility Pacific Gas and Electric Co. (PG&E) develop rebates for consolidation-driven data center deployments. The company's programs have become a model for other utilities in the U.S., and utilities in other parts of the world have also started to investigate rebates based on PG&E's model. Rebates are offered to data centers using virtualization technology to consolidate servers and produce energy savings. In addition to the rebate, PG&E estimates that customers can save \$300 to \$600 in annual energy costs for each server removed. These savings can almost double when reduced data center cooling costs are taken into account.

Research at Intel's own data centers

As you might expect, as the owner/operator of its own data centers, Intel also has a customer perspective on running more energy-efficient data centers. Over the past five years, the increasing power and cooling demands of servers and other IT equipment have challenged Intel® IT to develop strategies for rapidly increasing the power density its data centers can support.

To accommodate greater density, Intel IT has experimented with a holistic approach to designing air-cooled data centers that reduce the power not directly related to workloads. Intel has been able to achieve breakthrough power and heat densities of 15 kW per cabinet and more than 500 watts per square foot (WPSF) of server room area at a lower cost than most other data center designs. Based on extensive modeling, Intel believes the company can double that density in the future. Intel's approach includes several innovations:

- Modeling airflow to identify and address key airflow problems.
- Using barriers and custom cabinets to control airflow and increase the air conditioning airflow efficiency.
- Combining multiple techniques to achieve high densities in new and retrofitted buildings.

One successful technique Intel is using to lower power costs during three quarters of the year is a wet-side economizer. By exploiting evaporative cooling and satisfying a local building code requirement for energy conservation, Intel's latest data center in "economizer mode" achieves a 5 in the HVAC effectiveness index (ratio of IT equipment power to HVAC power), exceeding 22 other data centers studied. The best of these data centers could only achieve a 3.5 ratio and some of these data centers posted ratios under 1. This means they were actually using more power to cool the servers than run them.

What Intel's Efforts Mean for the Data Center of the Future

With power shortages and energy costs becoming growing concerns, a movement to improve data center and other IT energy-usage patterns is quickly gaining momentum as the pieces fall into place. For the present, data center operators can harness the benefits of Moore's Law by buying more powerful, energy-efficient servers based on Intel® multi-core processors and Intel Core micro-architecture to increase rack data center density. By consolidating operating systems and applications on these servers through virtualization, data center operators can reduce the number of physical servers required, thereby reducing floor space, cooling, and capital costs, and ultimately TCO.

In the not-too-distant future, data center operators can look forward to better metrics for understanding the overall energy efficiency of their data center in relation to compute output, as well as better benchmarks for comparing servers and other

equipment designed to help them improve their energy efficiency. Intel's and the industry's efforts to reduce power consumed and maximize compute performance at both the server and data center level will help the world handle its ever-increasing need for compute capacity.

Learn More

For more on Intel's efforts in energy efficiency, see:

www.intel.com/technology/eep/index.htm

For more on Intel's research into energy-efficient systems, see:

www.intel.com/research/platform/energy/index.htm

White Paper: Intel Is Leading the Way in Energy-Efficient Platforms:

download.intel.com/technology/eep/green-paper.pdf

White Paper: Increasing Data Center Density While Driving Down

Power and Cooling Costs: www.intel.com/business/bss/infrastructure/enterprise/power_thermal.pdf

www.intel.com/business/bss/infrastructure/enterprise/power_thermal.pdf

White Paper: Air-Cooled High-Performance Data Centers:

Case Studies and Best Methods: www.intel.com/it/pdf/air-cooled-data-centers.pdf

www.intel.com/it/pdf/air-cooled-data-centers.pdf

Visit The Green Grid: www.thegreengrid.org

Action list for turning your energy challenges into opportunities

- Recognize that the data center and the IT equipment are becoming an integrated system that needs to be thought of as a whole.
- Leverage industry organizations, such as Green Grid and ASHRAE for data center best practices.
- Remove the wall between IT and facilities by working together to improve cooling and power usage, distribution and conversion.
- Design data center facilities systems to use energy-efficient designs in power delivery, conversion and cooling.
- Demand and buy IT equipment that uses high-efficiency components.
- Understand data center metrics and optimize using the right ones.
- Take advantage of the Intel's continuing performance and energy-efficiency gains through Moore's Law.

¹Compared to 64-bit Intel® Xeon® processors at 110W TDP (110W per core).

²A. 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.

⁴Intel® processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See www.intel.com/products/processor_number for details.

[†]Intel® Virtualization Technology requires a computer system with an enabled Intel® processor, BIOS, virtual machine monitor (VMM) and, for some uses, certain platform software enabled for it. Functionality, performance or other benefits will vary depending on hardware and software configurations and may require a BIOS update. Software applications may not be compatible with all operating systems. Please check with your application vendor.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

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

Copyright © 2007 Intel Corporation. All rights reserved.

Intel, the Intel logo, Intel. Leap ahead., Intel. Leap ahead. logo, Intel Core, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

