

White Paper
Dual-Core Intel® Xeon® Processor
Data Center Optimization

Increasing Data Center Density While Driving Down Power and Cooling Costs

It takes a comprehensive strategy to scale data center capabilities, while simultaneously containing power and cooling costs. New Dual-Core Intel® Xeon® processor-based servers provide a critical new resource, by delivering leading performance, price/performance and energy-efficiency for a broad range of business applications. Read about this and other Intel advances that can help you increase density, reduce costs and extend the life expectancy of your existing facilities.

Executive Summary	ii
The High Cost of High Power	1
A Breakthrough in Energy-Efficient Performance	2
Optimizing Performance per Watt—System	4
Optimizing Performance per Watt—Data Center	5
The Critical Importance of Software	5
The Future of Energy-Efficient Performance	7
The Bigger Picture—Capabilities per Watt and TCO	8
Conclusion	10
Appendix A: Performance and Energy-Efficiency Tests and Comparisons	11
Appendix B: Best Practice Recommendations	12
Appendix C: Additional Resources	15
Sidebar: Lower TCO and Better Productivity for Enterprise Clients	4
Sidebar: Transforming the Data Center	5
Sidebar: Enthusiastic Support from Leading Server Vendors	6
Sidebar: Highest Density with Energy-Efficient Blades	9

Executive Summary

Relief has arrived for IT and facilities managers who need to pack more computing capacity into existing data centers, while simultaneously reducing total costs. For some time, Intel has been working to deliver new levels of energy-efficiency through silicon, processor, platform and software innovation. The results of these efforts are now clearly evident in the new generation of servers based on the Dual-Core Intel® Xeon® processor 5100 series (code-name Woodcrest). These servers boost performance by up to 3 times and energy-efficiency by more than 3 times compared to previous generation, single-core Intel Xeon processor-based servers. They are delivering the best performance, price/performance and energy-efficiency in their class, and can help IT organizations dramatically increase compute density, while reducing power and cooling requirements.

These processors also include silicon-level support for virtualization, another critical technology for optimizing data center power and cooling efficiency. By enabling multiple applications and operating systems (OSs) to be consolidated per server, virtualization can help IT organizations consolidate their server infrastructure, so there are fewer systems to power and cool—and more room for expanding compute capacity within existing facilities.

These and other recent innovations are major steps toward increasing performance and energy-efficiency, but they are only the beginning. Intel researchers continue to push the limits of transistor density in next-generation process technologies, to deliver better performance and new capabilities, while simultaneously driving down power consumption. Intel is also delivering software tools, training and support that help software vendors and corporate developers optimize their software for multi-core processors and 64-bit computing. These are essential efforts, since optimized software can boost performance, while containing or even reducing power consumption.

This paper discusses these and other advances that can help IT managers plan and implement a more comprehensive and effective strategy for expanding capabilities while reducing power consumption. It also includes a summary of data center best practices for optimizing the use of power and cooling resources, so IT managers can make better use of their current facilities, while reducing both capital and operational costs.

The High Cost of High Power

Data Center power and cooling challenges are increasing and are impacting IT and facilities managers on many levels.

- **High Utility Costs**¹—Power and cooling costs have become the second-largest contributor to data center total cost of ownership (TCO), second only to IT payroll. With the global trend toward rising energy rates, it seems almost certain that the impact on TCO will continue to grow.
- **Limited Compute Densities**—Today's high-density rack and blade servers help to reduce space, networking, cabling and management costs—as well as total power consumption. However, they drive up power density to levels that can exceed the limits of many facilities.²
- **Aging Data Centers**—As power and cooling needs continue to increase, they are causing many facilities to reach full capacity sooner than expected, and forcing companies to embark on expensive upgrades or new construction. Moreover, construction costs are increased by the need to address growing but unpredictable requirements.

These challenges would be relatively easy to address in a static IT environment. But businesses must constantly expand and modernize their infrastructure to stay competitive. As they do so, an effective power and cooling strategy is increasingly important to keep costs down and options open.

An Industry-Wide Challenge

“Making data centers more power-efficient will require radical improvements in power production by multiple parties, including semiconductors vendors, systems vendors and software vendors.”

— Richard Fichera, vice president, Forrester Research³

Barring the sudden emergence of a cheap and plentiful new energy source, there is no simple solution to today's power and cooling challenges. Addressing them will require a broad range of power reduction technologies, along with a comprehensive IT strategy that optimizes power usage across the data center. Key focus areas include:

- **Energy-efficient IT solutions**, including servers, storage, networking and communications systems, as well as power distribution and cooling systems that are more flexible, modular and efficient.
- **Better utilization of each system** through virtualization, consolidation and effective workload management.
- **Optimized software** that takes full advantage of new server capabilities (multi-core, 64-bits, virtualization, etc.) to deliver better performance with fewer and less energy-intensive compute cycles.
- **Best practices** across the data center to reduce utility costs and increase the capacity of existing facilities.

Intel is delivering value across all these areas today. The most dramatic advances are evident in the new generation of high-performance, energy-efficient Dual-Core Intel Xeon processors.

¹ According to Vernon Turner of IDC, the annual power and cooling bill for 100 servers is currently about \$40,000. Source: *Servers swamp data centers as chip vendors push ahead*, by Patrick Thibodeau, Computerworld, February 6, 2006: www.computerworld.com/hardwaretopics/hardware/server/story/0,10801,108433,00.html

² Although the power density of a fully-populated rack of blades is too high for most data centers, the lower power consumption per server can still deliver important benefits. According to APC, *“Contrary to popular belief, the primary TCO benefit related to Network-Critical Physical Infrastructure for blade servers comes from their reduced power consumption, NOT their reduced space consumption. Blade servers do not need to be installed at high density to obtain these TCO benefits.”* 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

³ Source: *Power and Cooling Heat Up The Data Center*, by Richard Fichera, Forrester Research, March 8, 2006. Available for purchase at: www.forrester.com/Research/Document/Excerpt/0,7211,38746,00.html

A Breakthrough in Energy-Efficient Performance


Dual-Core Intel Xeon Processor 5100 Series

The new Dual-Core Intel Xeon processor 5100 series (code-named Woodcrest) is fueling a major shift toward higher-performing and more energy-efficient servers. Based on the new Intel® Core™ microarchitecture, this processor boosts server performance by as much as 3 times compared to the fastest single-core Intel Xeon processors, while increasing energy-efficiency by more than 3 times.⁴ Systems based on these processors are now widely available from leading server vendors, and are delivering the best performance, price/performance and energy-efficiency in their class⁵ (Figure 1).

To help server vendors and IT organizations balance power versus performance in diverse environments, Intel is offering several processor versions, including 65W processors for mainstream servers; 40W processors for ultra dense implementations, such as server blades; and 80W processors for implementations where absolute performance is most important.

Intel Core microarchitecture builds on the power-saving technologies Intel originally implemented in the Intel® Pentium® M processor for laptops. Many of those technologies have now been extended and optimized for Intel server processors, and integrated with a number of additional capabilities that help to deliver high performance and low power consumption under heavy workloads.

Figure 1. New Dual-Core Intel Xeon processor-based servers deliver best-in-class performance and performance-per-watt across a broad range of applications and benchmarks, while also consuming less total power than competing solutions. This figure shows the results for the industry-standard SPECint_rate_base 2000* benchmark. For test parameters, system configurations, and a variety of additional benchmark and application test results, visit the Intel Web site at: www.intel.com/performance/server/xeon/index.htm



⁴ Based on Intel internal measurements made on May 3, 2006, comparing performance and power consumption with previous 64-bit, single-core Intel Xeon Processor based platforms. The new Dual-Core Intel Xeon processors also provide dramatic improvements versus previous generation dual-core Intel Xeon processors (up to 2 times better performance, and more than 2 times better energy efficiency). Comparisons to single-core systems are emphasized in this paper, because those are the measurements most relevant to customers who are considering replacing or upgrading older servers. 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.

⁵ Servers based on the Intel Xeon processor 5100 series have already set world records in 20 benchmarks. For more information, see the Intel press release: www.intel.com/pressroom/archive/releases/20060523corp.htm

Key advances include:

- **More instructions per clock-cycle**—Each core can process up to four simultaneous instructions, versus only three in previous generations. In addition, many common instruction pairs are combined into single instructions, to further improve processing efficiency.
- **Faster and more efficient data access**—A more efficient memory subsystem and a multi-core-optimized, shared-cache architecture significantly accelerate data access, so multiple cores can sustain higher levels of productivity.
- **Faster execution of key instructions**—The per-clock execution of Streaming SIMD Extension (SSE/SSE2/SSE3) instructions is effectively doubled, enabling a 4x increase in floating-point performance. This delivers critical speedups for a broad range of applications, including many security, financial, engineering and scientific solutions.
- **Dynamic power management**—Power delivery to key processor subsystems is dynamically managed during runtime, to continuously optimize performance versus power efficiency as workloads vary.

For more information, read the Intel white paper: *Inside Intel Core Microarchitecture: Setting New Standards for Energy-Efficient Performance*: ftp://download.intel.com/technology/architecture/new_architecture_06.pdf

Multi-Core Everywhere

Multi-core processors are one of the most effective strategies for multiplying compute capacity while maintaining or reducing power consumption. Intel's 65 nm process technology—combined with its global manufacturing capacity—is enabling very cost-effective, high-volume delivery of multi-core processors across all Intel's server and client platforms. As a result, dual-core processor-based systems will be widely available, and there will be no price premium for dual-core versus single-core processors. This enables IT organizations to standardize on dual-core processor-based systems today, so they can lay a foundation for high-performance, energy-efficient computing across both current and future needs.

Additional Energy-Efficiency at the Processor Level

Intel processors offer additional features that help to reduce power consumption in a wide range of hardware and software environments. The most important is Demand Based Switching (DBS), which 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.

Rapid and Ongoing Gains in Performance/Watt

“Every two years, we will bring out new silicon and new microarchitecture, instead of our [historical] four- to six-year cadence.”

– Paul Otellini, CEO, Intel Corporation⁶

Intel has a long-term processor roadmap for delivering ongoing advances, and more than 10 processors with four or more cores are in development today. In all cases, there is an intense focus on ramping performance and energy-efficiency through both per-core innovation and multi-core integration. This is an important point, since the majority of today's software is single-threaded and cannot take full advantage of multiple cores. By increasing both per-core and multi-core performance, Intel will provide an optimized, energy-efficient hardware platform for both current and emerging software applications.⁷

⁶ Source: *Intel steps up chip cadence*, by Stephen Shankland, CNET News.com, April 28, 2006: <http://news.zdnet.co.uk/0,39020330,39266019,00.htm>

⁷ According to ars technica, *“[Intel] Core [microarchitecture] looks like it has what it takes to carry Intel forward for at least another five years. By focusing on single-threaded performance, Core will excel on the types of applications that will make up the vast majority of server and consumer code in the near to medium term. And because it's designed for relatively low core-count multicore, it will help the software industry gradually make the transition to multithreaded code.”* Source: Into the Core: Intel's Next Generation Microarchitecture, by Jon “Hannibal” Stokes, ars technica, April 5, 2006: <http://arstechnica.com/articles/paedia/cpu/core.ars/1>

Optimizing Performance per Watt—System

Although energy-efficient processor performance is critical to overall server value, it is not enough. Without complementary advances in I/O and memory subsystems, bottlenecks can emerge that negate the advantages of new processors. To address this need, Intel focuses on all major server components, delivering comprehensive platform advances that help to optimize performance and energy-efficiency across diverse workloads.

Compared to previous generations, new Dual-Core Intel Xeon processor-based servers provide:⁸

- **Up to 3 times the processor interconnect bandwidth**
- **Up to 4 times the memory capacity**
- **Up to 3 times the peak memory bandwidth.** The memory subsystem (based on Fully Buffered DIMM Technology, see below) operates at full speed even at full capacity, unlike competitive offerings, which force IT organizations to make tradeoffs between speed and capacity.

The latest server platforms offer additional capabilities that help to address the performance and scalability demands of today's increasingly network- and data-intensive applications:

- **Intel® I/O Acceleration Technology (Intel® I/OAT)**—This technology optimizes network data processing across all critical platform components to increase I/O throughput by up to two times, with up to 40 percent less I/O related load on the processor.⁹ Both I/O and processing performance are improved with little or no increase in power consumption. Intel I/OAT requires no software changes and is designed to scale in future implementations.¹⁰
- **Fully Buffered DIMM (FBDIMM) technology with industry-standard DDR2-667 memory**—FBDIMM technology boosts memory capacity up to 4 times, and helps reduce latencies as memory loads increase. It also adds a number of memory reliability features to increase data integrity and resilience, and provides a

scalable architecture for future memory solutions.¹¹ Processor and I/O advances have been outpacing memory solutions for some time. Although FBDIMM technology adds some additional system-level power consumption, it brings the memory subsystem up to date, and will be increasingly critical to enable IT organizations to meet today's rapidly growing memory requirements simply, reliably and at low cost.¹²

- **Quad-Core Scalability**—Today's Dual-Core Intel Xeon Processor-based server platforms are also optimized for quad-core Intel Xeon processors. This will enable server vendors and IT organizations to support another major leap in performance and energy-efficiency using a common platform architecture that simplifies implementation and reduces total costs. (Quad core Intel Xeon processors were demonstrated in March 2006 at the Intel Developer Forum, and are expected in the first half of 2007.)

Lower TCO and Better Productivity for Enterprise Clients

The new Intel Core microarchitecture is delivering leading performance and energy-efficiency across all mainstream Intel computing platforms, including desktops and laptops. As with servers, Intel is complementing these new processors with comprehensive platform innovation that can dramatically improve flexibility, manageability and security. New Intel processor-based business PCs now include Intel® vPro™ technology, which helps to address some of today's most critical challenges, by helping IT organizations control, secure and maintain their client systems more efficiently than ever before. For more information, visit the Intel website at: www.intel.com/vpro/index.htm?iid=HMPAGE+Feature_06ww17

⁸ Comparisons are with respect to previous 64-bit Intel Xeon based servers available in Q2'05. The full 3x peak memory bandwidth improvement is based on using 667MHz FBD Memory technology.

⁹ Intel internal measurements based on comparisons with respect to previous 64-bit Intel Xeon processor-based servers available in Q2'05. Intel I/O Acceleration Technology requires use of Dual-Core Intel Xeon Processor 5000 Sequence Processors, Intel® 5000 Sequence Chipsets, Intel® 6321 ESB I/O Controller Hub, either Intel® 82563EB/82564EB or Intel® PCIe Server Adapter with Intel's Nyssa 4.1 Beta Release or later, Microsoft Server 2003 with Scalable Network Pack or Linux 2.6.12 Kernel (or later).

¹⁰ For more information about Intel I/O Acceleration Technology and its benefits, see: www.intel.com/technology/magazine/communications/intel-ioat-0305.htm

¹¹ According to *The Inquirer*, "...it (FBDIMM) will make the transitions to completely new memory architectures vastly easier and probably quicker. It effectively decouples the logic of the memory controller from the memory architecture without adding much delay. This is a good thing. Let me repeat, a very good thing." Source: Intel FB-DIMMs to offer real memory breakthroughs, by Charlie Demerjian, *The Inquirer*, April 5, 2004: www.theinquirer.net/?article=15167

¹² For more information on FBDIMM, see the Intel white paper, *Fully-Buffered DIMM Technology Moves Enterprise Platforms to the Next Level*: www.intel.com/technology/magazine/computing/Fully-buffered-DIMM-0305.htm

Optimizing Performance per Watt—Data Center

Virtualization and Consolidation

Virtualization is a fundamental technology for reducing power and cooling requirements, as well as for general data center optimization. With today's virtualization software, IT organizations can consolidate multiple OS and application stacks on each server. This reduces total power and cooling requirements by increasing utilization and allowing fewer servers to handle a given workload. It also helps to improve manageability, availability, security and agility.

With their dual-core processors, large memory capacity, high-bandwidth I/O and leading RAS capabilities, the latest Intel Xeon processor-based servers are ideal for virtualized server environments. They also include Intel® Virtualization Technology, which improves the interoperability, robustness and supportability of today's software-based solutions, such as VMware ESX Server*, Microsoft Virtual Server* and Xen*.

By supporting core virtualization processes in silicon, this technology enables important new capabilities, such as mixed guest OS environments (including both 32-bit and 64-bit versions of Windows and Linux). It will also help to speed innovation, by simplifying virtualization software solutions.

Intel Virtualization Technology is the first step in a multigenerational roadmap of increasingly powerful virtualization enhancements. Intel is working with key virtualization software vendors and a broad base of application software vendors to deliver best-in-class virtualization capabilities that help to improve ease of use, value and reliability. Intel's unique ability to collaborate across the virtualization ecosystem is critical in developing optimized and broadly supported virtualization technologies that support new capabilities and greater value on proven, industry-standard platforms that can be deployed with confidence in diverse environments. For more information, visit the Intel Web site at: www.intel.com/business/bss/products/server/virtualization.htm

Transforming the Data Center

- 75 percent fewer racks
- 63 percent lower TCO

Energy-efficient servers mean more than reduced power and cooling costs. They also mean higher compute density, which translates to fewer servers and fewer racks for a given workload. According to Intel estimates, new Dual-Core Intel Xeon processor-based servers can reduce the number of racks needed to support a given workload by up to 75 percent for a given workload, and save up to 63 percent in total data center costs (compared to previous-generation, single-core Intel Xeon processor-based servers).^a

This is a major advantage in any data center environment, and particularly for high-performance computing (HPC) implementations, where density and energy-efficiency are of critical importance. It is also worth noting that the above savings are based on average power, cooling and space costs in a typical data center. They do not include potential savings from postponing data center upgrades and new construction that might otherwise be necessary.

^a Based on typical data center power and cooling costs, plus \$10/sqft for system maintenance and administration. Performance estimates are based on the industry-standard SPECint_rate benchmark, and would vary depending on specific systems, configurations, workloads and IT environments. The highest densities and cost savings are realized with low-voltage versions of new Dual-Core Intel Xeon processors.

The Critical Importance of Software

Multi-core processors deliver immediate performance and energy-efficiency benefits for multi-threaded applications, which include all major OSs 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.

Enthusiastic Support from Leading Server Vendors

Dell

"Dell believes in delivering scalable, high-performance computing solutions in high volumes to enable easy and highly cost-effective deployment and management. The latest Dual-Core Intel Xeon processors help us do that more effectively than ever before. With the new levels of performance and energy-efficiency that Intel's new Core Microarchitecture brings—and leading redundancy and virtualization capabilities—they're helping our customers take data center efficiency to new levels."

— Neil Hand, Vice President of Worldwide Enterprise Marketing, Dell Product Group

HP

"The upcoming portfolio of industry-leading HP ProLiant 2P servers combines new Dual-Core Intel Xeon processors with HP's balanced system architecture to offer customers cutting-edge technologies that deliver leading performance and performance per watt. By advancing the key subsystems of the new 2P ProLiant servers, which are optimized for virtualization and leverage the leading Infrastructure Management software, customers are able to take full advantage of the performance to grow and align their IT infrastructure with changing business needs."

— Paul Miller, Vice President, HP Industry Standard Servers

IBM

"With IBM's mainframe-inspired Xtended Design Architecture (XDA), IBM System x servers with Dual-Core Intel Xeon processors deliver dramatic new performance capabilities with significant power savings built into every server. IBM System x servers introduced in May 2006 include the second generation of XDA, featuring scalable memory and I/O design to support business growth; virtualization optimization to ensure maximum resource utilization; and the IBM's PowerExecutive infrastructure based power management solution to improve power utilization and reduce energy costs."

— Alex Yost, Director, IBM System x

However, the majority of today's applications are single-threaded, so they cannot take direct advantage of multiple cores. Optimizing software for multi-threaded throughput can accelerate performance and improve utilization, which reduces the amount of energy required to accomplish the same work. 64-bit migration can deliver similar benefits for data-intensive applications, enabling faster response times and better scalability within approximately the same power envelope.

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

- **World-Class Software Development Tools**—Intel compilers, performance analyzers, threading tools and libraries are now used by over 200,000 commercial, open source and corporate developers. These tools are optimized for multi-threaded throughput and are compatible with popular development environments. They help to deliver leading performance and energy-efficiency across all compatible server platforms, while delivering best results for Intel processor-based systems (because they are highly tuned to take advantage of Intel's unique architectural features).
- **Optimized platform software**—Intel develops and optimizes much of the core platform software, so basic functions are quicker and less energy-intensive.
- **Deep collaboration with leading software companies**—Intel works with top OS, database and application vendors to further improve performance and energy-efficiency on the latest multi-core Intel processor-based platforms.
- **Extensive training and support**—The Intel® Software Network (www.intel.com/cd/ids/developer/asmo-na/eng/index.htm) and the Intel® Software College (<http://or1.cedar.cps.intel.com/softwarecollege/>) offer training, information and resources that can help software developers optimize their software more effectively and at less cost.

The Future of Energy-Efficient Performance

Intel researchers are paving the way to an estimated 10 times improvement in both performance and energy efficiency over the next 10 years, through ongoing advances in silicon, processor and platform technologies (Figure 2).¹³

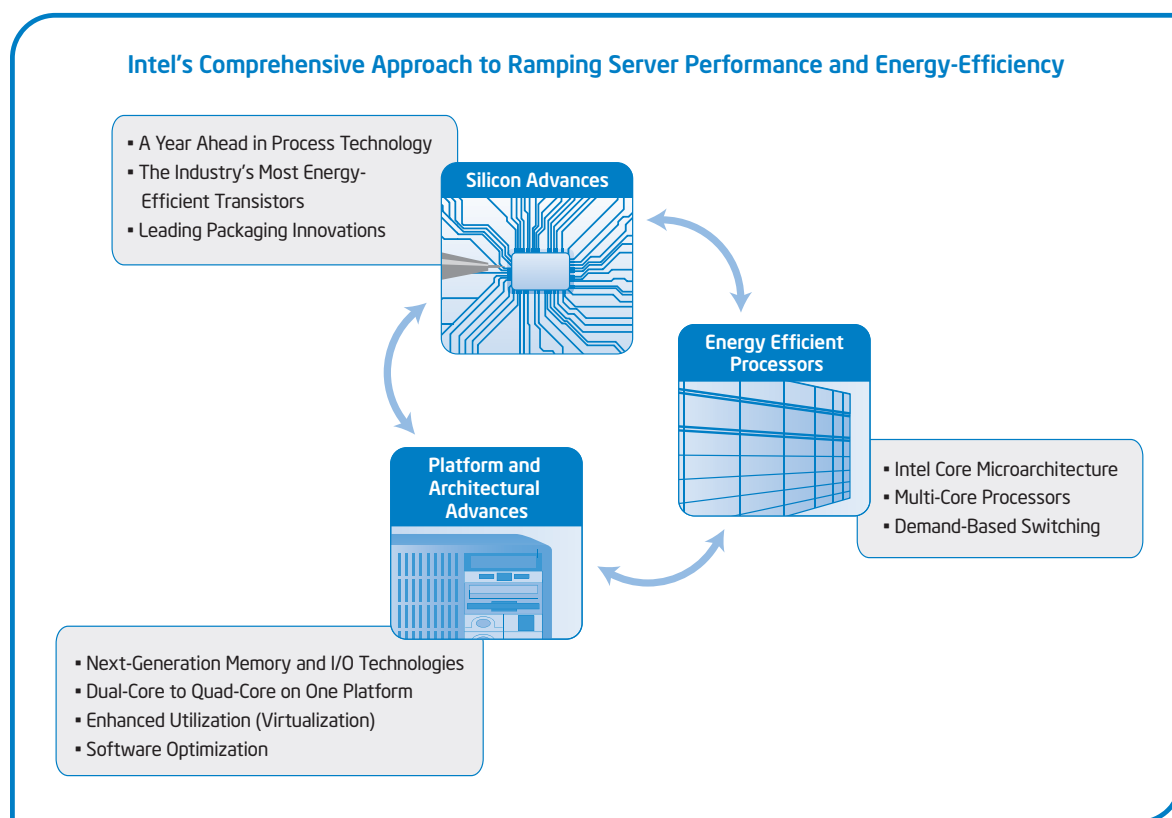
A Year Ahead in Silicon Process Technology

Intel is a full year ahead of the industry in silicon process technology, and has developed the most energy-efficient transistors for both 90 nm and 65 nm implementations.¹⁴ The company is currently moving to 65 nm manufacturing across all volume CPU products, and the full crossover for processors is expected in the third quarter of 2006.

Intel is also on track to begin high-volume manufacturing using its new 45 nm process technology in 2007. This technology is expected to double transistor densities once again, and has already been proven in an SRAM pilot project.¹⁵ It will give Intel design engineers the option of improving transistor switching speeds by up to 20 percent, or reducing leakage power by a factor of 5. It will also deliver as much as 30 percent improvement in transistor switching power.¹⁶ These advances provide a clear path to ongoing performance and energy-efficiency gains.

Intel has been doubling transistor densities about every two years for more than 40 years, and a number of recent breakthroughs provide a strong foundation for continuing this rate of innovation.

Figure 2. Intel is taking a comprehensive approach to solving today's power and thermal challenges, from silicon, platform and software innovations that will dramatically increase server efficiency, to system- and data center-level initiatives that involve broad industry collaboration.



¹³ "Intel can talk not only about a road map of two cores, leading to four cores, leading to eight and so on, but also the manufacturing processes that will produce line widths of 65 nanometers, dropping to 45 nanometers, dropping to 32 nanometers, dropping to 22 nanometers. And, more to the point, company engineers can tell you what year they'll be able to achieve these milestones." Roger L. Kay, president, Endpoint Technologies Associates Inc. Source: *How Intel Keeps Its Enterprise Customers Coming Back for More*, eWeek, March 10, 2006: www.eWeek.com/article2/0,1895,1936680,00.asp

¹⁴ Based on the published results of competitors.

¹⁵ SRAM (Static Random Access Memory) devices are commonly used to demonstrate technology performance, process yield and chip reliability prior to high-volume manufacturing of processors and other logic chips using a new process technology.

¹⁶ For more information, see: www.intel.com/technology/silicon/new_45nm_silicon.htm

As one example, Intel and QinetiQ have successfully produced an 85 nm transistor using indium antimonide (InSb). Tests show this new material can be used to reduce power consumption by up to a factor of 10 and improve performance by up to 50 percent compared to current transistors. With this and many other promising avenues of research, Intel expects to continue historical rates of advance through at least 2020.

For more information, visit the Intel Web site at: www.intel.com/technology/eep/index.htm

Multi-core and Beyond

Many Intel researchers are focused on finding the best ways to use the many billions or even trillions of energy-efficient transistors that will be available for future chip designs. Intel processors are likely to include dozens or even hundreds of general-purpose cores, along with specialized cores that deliver major performance and energy-efficiency benefits for particular functions, such as XML processing (XML traffic is on track to exceed all Web traffic by 2007). Intel is also exploring the possibilities for integrated workload managers that will allocate processor resources in the most effective and

energy-efficient manner based on specific workloads. These and many other strategies can be expected to dramatically increase both absolute performance and performance per watt in future Intel processors.

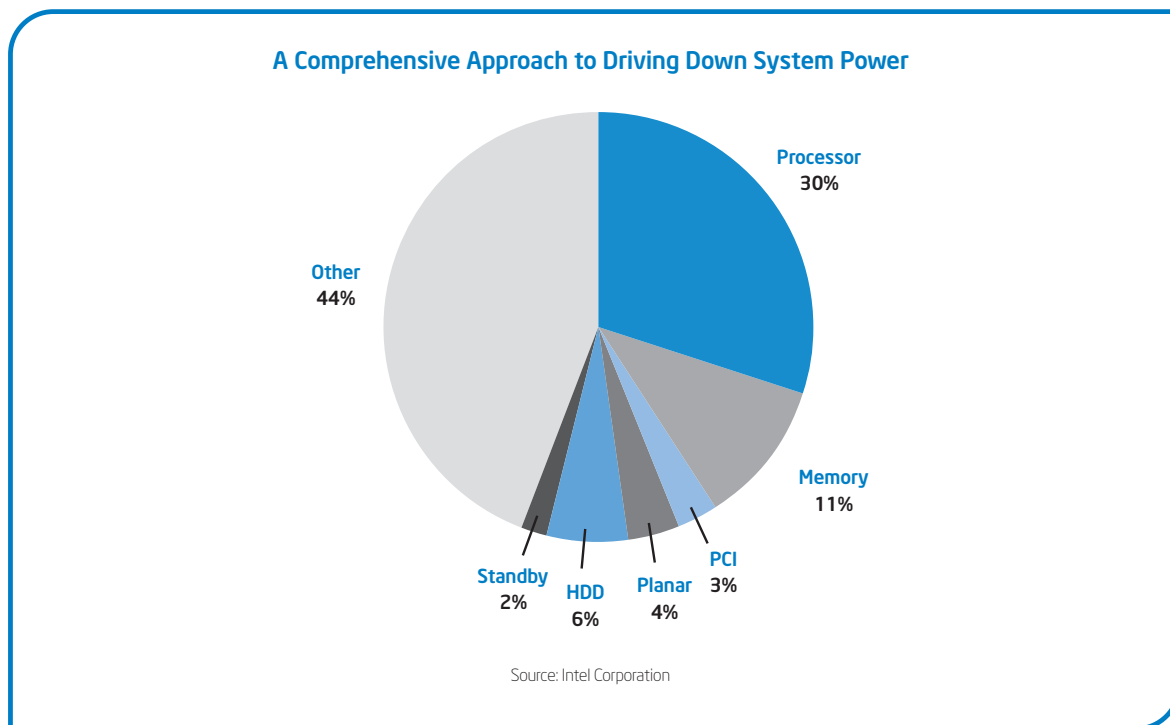
With the increased energy-efficiency of today's Intel processors, other platform components are contributing a higher percentage of the total power and cooling burden (Figure 3). Intel is working to increase energy-efficiency across the entire platform, and is leading the way toward system power reductions of 80 percent or more. This is an important benefit of Intel's comprehensive platform approach to server design, enabling optimization of energy-efficiency across all critical hardware components.

The Bigger Picture—Capabilities per Watt and TCO

“While the environmental and energy conservation aspects represent positive goals, the real issue is reducing the first-cost and lifecycle costs of the data center.”

— Michael Bell, research vice president, Gartner, Inc.¹⁷

Figure 3. Server processors account for just one component of server power consumption. Intel is working to increase energy-efficiency across the entire platform, and is leading the way toward system power reductions of 80 percent or more.



¹⁷ Source: *Energy Consortium Takes Aim at Bottom Line*, by Shamus McGillicuddy, SearchCIO.com, May 1, 2006: http://searchcio.techtarget.com/originalContent/0,289142,sid19_gc1186068,00.html?track=NL-48&ad=545073&Offer=52HTMLA

Although increasingly important, optimizing energy-efficiency is only one factor in the larger goal of driving business success through affordable and effective IT. Manageability, security, reliability, interoperability and standards compliance are all important in providing a flexible and affordable foundation for growth.

To deliver this broader value, Intel looks beyond performance-per-watt in designing next-generation server platforms. The larger goal is to optimize total platform capabilities per watt and per dollar, so that servers not only deliver better performance and energy-efficiency, but are also easier and less expensive to deploy, provision, manage and protect. Intel integrates leading technologies across all these needs. Many have been discussed already in this paper. Other key examples include:

- Leading RAS features**—The latest new RAS (reliability, availability and serviceability) features in Dual-Core Intel Xeon processor-based servers include Enhanced ECC (Error Correcting Code) and Memory CRC (Cyclic Redundancy Check), which improve the ability of the system to detect and correct memory errors.¹⁸ Other additions include support for memory mirroring and memory sparing. These advanced RAS capabilities not only improve the reliability of the memory subsystem, but also help to improve the integrity of hosted data. These and many other built-in RAS features deliver critical benefits, especially in today's virtualized server environments, in which system downtime can impact multiple, business-critical applications.
- More efficient management**—Automating management is key to improving IT efficiency and driving down TCO. Intel® Active Server Manager builds on Intel's recent breakthroughs in client manageability (via Intel® Active Management Technology). It enables more effective and resilient remote administration and management, including discovery, inventory, monitoring and recovery. It can allow centralized personnel to manage more systems with greater efficiency, and to accelerate time to repair in a wide range of common scenarios, including power failures and crashed OSs. For more information, visit the Intel Web site at: www.intel.com/cd/channel/reseller/asm-na/eng/250641.htm

Highest Density with Energy-Efficient Blades

"...a dual-core blade with the industry's best performance per watt..."

— IBM press release, February 8, 2006, referring to a blade based on the Dual-Core Intel Xeon processor LV 2.0GHz

Blade servers offer the ultimate value in compute density, manageability and IT agility, but adoption has been slowed by the high power and cooling requirements of these ultra-dense systems. Low-voltage versions of the latest dual-core Intel processors offer exceptional advantages for blade systems, and can help IT organizations achieve unprecedented levels of compute density with less draw on data center power and cooling resources.

Intel currently offers Dual-Core Intel Xeon processors that consume as little as 31W (for 32-bit applications only) and 40W (for both 32-bit and 64-bit workloads). As companies virtualize and consolidate their server infrastructure, they should give careful consideration to blade servers for optimizing the efficiency and flexibility of their data centers.

Most important, all Intel server technologies are developed, integrated and validated together, and delivered on the industry's most widely deployed server architecture. The result is platforms that deliver better total value, and can be deployed with confidence across the widest range of hardware and software environments. IT organizations can count on smooth and comprehensive platform improvements with every new server deployment and upgrade. They not only get leading performance and energy efficiency, but also better value across a broad range of IT requirements—all on a consistent, standards-based architecture that helps them streamline operations and reduce TCO.

¹⁸ Enhanced ECC enables retransmission of data not only in the event of a single-bit error, but also in the event of a double-bit error. With Memory CRC, address and command transmissions are automatically retried if a transient error occurs. This helps to prevent silent data corruption, which has the potential to bring down a system and/or lead to errors in business information (Memory CRC is provided only in systems configured with Fully Buffered DIMM memory technology).

Conclusion

New Dual-Core Intel Xeon processor-based servers provide breakthrough value in performance and energy-efficiency, while delivering a number of critical new technologies that improve headroom, flexibility, reliability and manageability across a wide range of mainstream hardware and software environments.

These new server platforms are the first step in a comprehensive drive to ramp server performance and density, while helping businesses drive down their data center power and cooling requirements. Through comprehensive platform innovation, software optimization and next-generation silicon process technologies, Intel will continue to help businesses expand their IT capabilities in the most cost-effective and energy-efficient manner, so they can reduce their total costs and extend the life and value of their current and future data centers.

Appendix A: Performance and Energy-Efficiency Tests and Comparisons

New Intel Xeon processor-based servers are demonstrating best-in-class performance, price/performance and energy-efficiency across a wide range of industry benchmarks and business applications. See the following documents and Web sites for specific test results.

- **General Server Performance**—This Web site documents performance and performance per watt leadership for numerous industry benchmarks, as well as for many business and HPC applications. It includes head-to-head comparisons with competing x86 servers:
www.intel.com/performance/server/xeon/index.htm
Many of these results are also summarized in an Intel press release:
www.intel.com/pressroom/archive/releases/20060523corp.htm
- **Workstation Performance**—*"Intel's Woodcrest processor previewed,"* by Scott Wasson, The Tech Report, May 23, 2006—This article provides a detailed look at Intel's new processor architecture, and includes a variety of performance and performance-per-watt comparisons based on popular workstation applications: www.techreport.com/etc/2006q2/woodcrest/index.x?pg=1
- **Financial Application Performance**—*"64-bit financial application-based workload performance on Intel- and AMD-processor-based server platforms."* This report provides detailed performance comparisons for two leading financial services applications (SunGard Credit Risk and the Black-Scholes Kernel): www.techreport.com/etc/2006q2/woodcrest/index.x?pg=1

Appendix B: Best Practice Recommendations

The following recommendations come from Intel® Solution Services and Intel's internal IT organization. Please be aware that specific cooling challenges can be complex, and many organizations will benefit from consulting with outside specialists.

Platform

1. Take Advantage of New Energy-Efficient Dual-Core Intel Xeon Processor-Based Systems.

They can improve performance by up to 3 times, while reducing power consumption (and heat generation) by more than 3 times.

2. Consolidate Workloads to Boost Utilization—Multiple OS and application stacks can be consolidated on fewer servers using virtualization software, such as VMware ESX Server, Microsoft Virtual Server or Xen. This can substantially reduce total power consumption by making better use of existing infrastructure. Aggressive workload management can also help to increase data center power efficiency.¹⁹

3. Configure New Servers with DDR2 Memory—It has been shown to increase performance by up to 11 percent for memory-intensive workloads, while decreasing memory power consumption by 30 to 40 percent.²⁰

4. Move to Newer OSs for All Platforms—Microsoft Windows* Server 2003 supports Demand Based Switching (DBS), which can substantially reduce average system power consumption for servers operating at typical data center utilization rates. It also makes use of "halt states" that reduce power consumption even on older platforms that do not support DBS.

Rack

Most system-, rack-, and room-level cooling issues are created due to insufficient airflow or inadvertent mixing of hot and cold air. The need for sufficient airflow is obvious, but is often overlooked by IT personnel focused on other concerns. The mixing of hot and cold air is a more subtle issue, but equally problematic since it causes hot air to bypass the cooling infrastructure, thereby reducing cooling efficiency and often creating additional airflow issues. (To quantify cooling efficiency, Intel IT measures Air Conditioning Airflow Efficiency (ACAIE) in its facilities, which is simply the amount of heat that can be removed per standard cubic foot per minute of cooling air. As bypass air is reduced, ACAIE increases.)

1. Understand Airflow Requirements for Specific Equipment—There are four basic airflow scenarios: front-to-back, side-to-side, bottom-to-top, and top-to-bottom. Understanding the requirements for specific equipment will enable an efficient rack-level design and cooling strategy.

¹⁹ According to Forrester Research, "...increasing the average utilization of the server can yield significant benefits in overall operational efficiency. You can increase utilization by consolidating workloads—which has become substantially easier with system virtualization technology—or by managing workloads more actively." Source: *Power and Cooling Heat Up The Data Center*, by Richard Fichera, Forrester Research, March 8, 2006. Available for purchase at: www.forrester.com/Research/Document/Excerpt/0,7211,38746,00.html

²⁰ Based on internal Intel tests comparing performance and power consumption for an Intel Xeon processor-based platform with an Intel E7520 chipset and an 800 MHz FSB. Tests were run with DDR2-400 memory versus the same platform with DDR-333 memory. Actual performance and power savings will depend on specific workloads and server configurations.

2. Standardize on Racks Designed for High-Density Environments—Standardizing on an appropriate rack design makes it much easier to establish and enforce effective power and thermal policies. Avoid shallow racks to ensure in-rack cabling does not obstruct airflow. Consider racks that support retrofit fan or cooling units (but verify the benefits of these add-on units). The ability to supplement cooling for individual racks may be important to accommodate high-density systems without compromising roomwide efficiency.²¹

3. Arrange Racks in Rows to Establish Hot and Cold Aisles—Racks should be aligned front-to-front along cold aisles, and back-to-back along hot aisles. Within each row, racks should be tightly abutted. For this strategy to be effective, cold air must be delivered to cold aisles and hot air extracted from hot aisles. Work to eliminate hot-air remixing, which will cause short cycling of the cooling system.

4. Use Blanking Panels and End Baffles—Blanking panels improve airflow through the rack, minimize air loss, and help to prevent exhaust air recirculation. End baffles perform a similar function, and provide similar benefits, at the ends of data center rows.

5. Ensure Adequate Airflow to Individual Racks and Systems—Clearly define power and cooling requirements at the room, row, and cabinet level. Ensure sufficient airflow to racks based on system-level inlet air temperature and airflow requirements, and use thermal and aerodynamic analysis tools to model and design your cooling solutions. Insufficient airflow will often result in hotter systems and turbulence that decrease cooling efficiency. For example, if a rack requires more cold air than the room provides, its fans will pull in a mix of hot and cold air. This will result in reheating of the room, hotter systems, unhealthy airflows, and a substantial reduction in the efficiency of the cooling system.

6. Explore the Benefits of Blade Servers—Blade architectures may reduce total power consumption (per unit of compute power) and deliver substantial TCO benefits through reduced cabling, easier provisioning, and improved modularity. However, they may increase power and cooling density. It is therefore important to look at total costs, risks and benefits within your particular physical and operational environment. As you evaluate solutions, be aware that fully-populated racks are rarely necessary to realize the primary TCO benefits, which are more dependent on the reduced power and cooling requirements per server.

Data Center

1. Understand Data Center Airflow—The locations of cooling systems and ductwork are obviously critical, but so are the locations of racks, cable trays, firewalls, and other infrastructure elements. Blank off any floor opening that allows access air to escape the plenum. Software tools are now available that greatly simplify airflow and thermal analysis. Consider consulting with cooling specialists for complex implementations.

2. Optimize Room Temperature Settings—Consider increasing the Delta T of your cooling system to more closely match IT equipment specifications. This may allow you to reduce total airflow, while meeting the same cooling capacity and reducing operational costs. (As an example, Intel IT has found it beneficial to lower supply air temperatures to between 55 and 65 degrees Fahrenheit while increasing Delta T values to 26 degrees Fahrenheit.)

²¹ According to Forrester Research, “Several sources report that local cooling and/or closed loop rack cooling can save in excess of 20 percent over cooling with conventional whole-room systems.” Source: *Power and Cooling Heat Up The Data Center*, by Richard Fichera, Forrester Research, March 8, 2006. Available for purchase at: www.forrester.com/Research/Document/Excerpt/0,7211,38746,00.html

- 3. Pay Attention to Infrastructure Efficiency**—It is generally worthwhile to spend more for infrastructure components that run efficiently at anticipated loads. Power loss in uninterruptible power supplies, power distribution units, cooling systems, etc., just add to the thermal load.
- 4. Perform Regular Power and Thermal Audits**—New systems, upgrades, and room changes can have unintended consequences, so it is important to monitor airflow, temperature, and other environmental factors regularly.
- 5. Avoid Over-Design**—Right-sizing power and cooling infrastructure—and data centers—is one of the most effective ways to reduce capital and operational costs in the data center. Work to understand life-cycle requirements and size infrastructure accordingly. Track vendor innovations, and, whenever possible, move toward more modular, flexible, and standardized solutions that improve agility and scalability.²²
- 6. Establish Policies and Educate Personnel**—Best practices for power and thermal management must become an integral component of data center operations. Everything from temperature and humidity settings to new system and cable deployments should follow well-understood guidelines that optimize cooling efficiency and minimize airflow obstructions and hot/cold air mixing.
- 7. For New Data Centers, Establish a Master Plan Based on Usage**—Different usage models require different layouts and capacities to enable optimized cooling solutions (Table B1). In designing and upgrading facilities, be aware of the trade-offs between system density and data center space. Higher densities can create design challenges, however, spreading out systems requires longer conductor runs for the power infrastructure. It also increases the total space that must be cooled (requiring larger chillers and additional fan capacity). Intel data centers for design teams currently target approximately 525 watts/sq foot (up to 15kW per rack). Airflow modeling shows that value can be realistically increased up to 1 200 watts/sq foot.

Table B1. Understanding Power Requirements for New Data Center Designs

	Data Center Focus	
	Technical Design	Business Transaction
Typical Configuration	60% High Density 30% Low Density 10% Network Devices	30% High Density 50% Low Density 20% Network Devices
Typical Rack Power	High Density (recent past): 3kW-10kW High Density (current): 10kW-14kW Blade Racks (current): 15kW-20kW	Rack load steady at 3kW-10kW (High-density systems can be spread out to dissipate heat).

²² According to Forrester Research, “A modular approach to provisioning and cooling the incremental space as it is needed can pay off with major gains in efficiency. According to several HVAC experts, one of the major defects in many data centers is that they are just too big, and if made smaller, the inventory could be cooled much more efficiently.” Source: Power and Cooling Heat Up The Data Center, by Richard Fichera, Forrester Research, March 8, 2006. Available for purchase at: www.forrester.com/Research/Document/Excerpt/0,7211,38746,00.html

Appendix C: Additional Resources

Intel Power & Thermal Web Resources

Addressing power and thermal challenges is an ongoing effort that encompasses a wide range of technologies and strategies. Get the latest information on related Intel solutions and research at: www.intel.com/technology/eep/

Intel® Solution Services

As businesses grow their IT infrastructure, they must occasionally make critical decisions about buying, building, optimizing, or consolidating data center facilities—decisions that can impact business and IT options for many years. Effective solutions must take into account not only changing business and IT requirements, but also technology trends that may substantially alter needs and capabilities over the life of the facility or equipment. Intel Solution Services, the worldwide consulting organization within Intel, offers expert technical assistance and consultation based on formal, best-known methods for identifying ways to optimize the productivity of current facilities, projecting future usage models and capacity requirements, and developing long-term strategic plans that balance business opportunity against TCO. Intel Solution Services has worked successfully with many clients to develop detailed business, technical, and financial analyzes that drive smarter and faster decisions on crucial data center issues. For more information, visit: www.intel.com/go/intelsolutionservices

IT@Intel: Discover How Intel Does IT

Intel's IT organization is a critical component in every aspect of Intel's business. It runs a global infrastructure consisting of more than 80,000 servers and 140 data centers (450,000 square feet) with more than 50 MW of power and cooling worldwide. It also provides products and services to Intel's external customers, including help desk services to 80,000 individual IT users and information solutions to Intel Business Groups. The IT@Intel program offers a variety of white papers and other resources providing useful insights into critical IT issues. For more information about Intel IT, visit: www.intel.com/it/

Other Data Center Management Resources

7x24 Exchange: www.7x24exchange.org/

AFCOM: www.afcom.com

APC: www.apc.com

Data Center Energy Management: <http://hightech.lbl.gov/DCTraining/top.html>

Dell: Assessing Power and Cooling Requirements White Paper:

www.dell.com/downloads/global/power/ps3q05-20050115-Moss.pdf

Data Center Environment Assessment:

www.dell.com/content/topics/global.aspx/services/en/dps_dcea?c=us&cs=555&l=en&s=biz


Forrester: www.forrester.com

Gartner: www.gartner.com/

HP Mission Critical Support: www.hp.com/hps/mission/

Microsoft: www.microsoft.com/management/default.mspix

Rackable Systems: www.rackable.com/index.htm



www.intel.com/go/xeon

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. Intel products are not intended for use in medical, life saving, life sustaining, critical control or safety systems, or in nuclear facility applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Intel 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, the Intel logo, Intel. Leap ahead and Intel. Leap ahead logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others. Information regarding third-party products is provided solely for educational purposes. Intel is not responsible for the performance or support of third-party products and does not make any representations or warranties whatsoever regarding quality, reliability, functionality, or compatibility of these devices or products.

Copyright © 2006 Intel Corporation. All Rights Reserved.

0606/JR/IMP/XX

313462-001US

