White Paper

Dual-Core Intel® Xeon® Processor

High Performance Computing

Solving Power and Cooling Challenges for High Performance Computing

Recent breakthroughs target some of today's toughest requirements

It takes a comprehensive strategy to scale high performance computing (HPC) capabilities, while simultaneously containing power and cooling costs. New Dual-Core Intel® Xeon® and Intel® Itanium® processor-based servers offer a critical new resource, delivering dramatic increases in performance, price/performance and energy-efficiency across a broad range of HPC applications. Read about this and other Intel advances that can help you increase density, reduce costs and scale capacity in your existing facilities.

June 2006

Executive Summary

Relief has arrived for organizations that need to pack more computing capacity into existing high performance computing (HPC) facilities, while simultaneously reducing power and cooling costs. For some time, Intel has been focused on helping IT managers address these issues, by driving new levels of energy-efficiency through silicon, processor, platform and software innovation. The results of these efforts are clearly evident in the new Dual-Core Intel® Xeon® processor 5100 series (code-name Woodcrest) and the upcoming Dual-Core Intel® Itanium® 2 processor 9000 series (code-name Montecito), which dramatically increase performance and energy-efficiency compared to previous generations.

These and other recent innovations are major steps toward increasing density, pure performance, price/performance and energy-efficiency for HPC solutions, but they are only the beginning. Intel researchers continue to push the limits of transistor density in next-generation process technologies, while simultaneously driving down power consumption. Intel is also delivering software tools, training and support that help developers optimize their software for multi-core processors and 64-bit computing. These are essential efforts, since optimized software can substantially boost performance and system utilization, while contributing to the containment or even reduction of 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 HPC capabilities, while improving density and reducing power and cooling costs.

The High Cost of High Power

"Making data centers more power-efficient will require radical improvements in power production by multiple parties, including semiconductor vendors, systems vendors and software vendors."[1]

– Richard Fichera, vice president, Forrester Research

HPC power and cooling challenges are increasing, impacting IT and facilities managers on many levels.

- **High Utility Costs**[2]—Power and cooling costs have become the second largest contributor to 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.[3]
- **Aging HPC Facilities**—Rising power and cooling needs are causing many facilities to reach full capacity sooner then expected, forcing organizations to embark on expensive infrastructure 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 HPC managers must constantly expand and modernize their HPC capabilities to keep pace in their industry or field of research. As they do so, an effective power and cooling strategy is increasingly important to keep costs down and options open.

An Industry-Wide Challenge

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:

- **Dense and energy-efficient building blocks,** including servers, interconnect and storage systems, as well as power distribution and cooling solutions that are more flexible, modular and efficient.

Higher Density Means Better Value for HPC Solutions

- Up to 75 percent fewer racks
- Up to 63 percent lower TCO
- Extended lifecycle for existing facilities

Energy-efficient servers mean more than reduced power and cooling costs. They also mean higher compute density, which translates to more servers in fewer racks—and more compute capacity within existing facilities. According to Intel estimates, new Dual-Core Intel Xeon processor-based servers can reduce rack requirements by as much as 65 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 infrastructure 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.

---

[1] 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

[2] 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

[3] While the power density of a fully populated rack of blades is too high for most facilities, 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

- **Better utilization of each system** through optimized solution and software design. Optimizing solution design (server configurations, cluster sizes, interconnect technologies, etc.) is critical to achieve maximum compute capacity per system. Software optimization is also important to take full advantage of the latest server technologies (multi-core, 64-bits, etc.). Relatively small efforts in optimization can result in substantial performance improvements within the same power and thermal envelope.
- **Optimized power and cooling infrastructure** based on actual systems and workloads. Specifications from server vendors typically assume maximum workloads, which are rarely the case in HPC environments. It is therefore important to test requirements under actual application workloads to avoid costly over-design.
- **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 processors.

A Breakthrough in Energy-Efficient Performance

Dual-Core Intel® Xeon® Processor 5100 Series

"Intel's newest Xeon has taken back the performance/Watt crown. In one word: Woodcrest rocks!"
– Johan De Gelas, June 7, 2006[4]

The new Dual-Core Intel Xeon processor 5100 series (code-name 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 with the fastest single-core Intel Xeon processors, while increasing energy-efficiency by more than 3 times.[5] Servers based on these new processors are widely available from leading server vendors, and are delivering the best performance per watt in their class across a wide range of benchmarks (Figure 1), along with leading absolute performance and price/performance.[6]

Figure 1. New Dual-Core Intel Xeon processor-based servers offer breakthrough performance per watt (as well as leading absolute performance and price/performance), all on the industry's most flexible and widely deployed architecture.

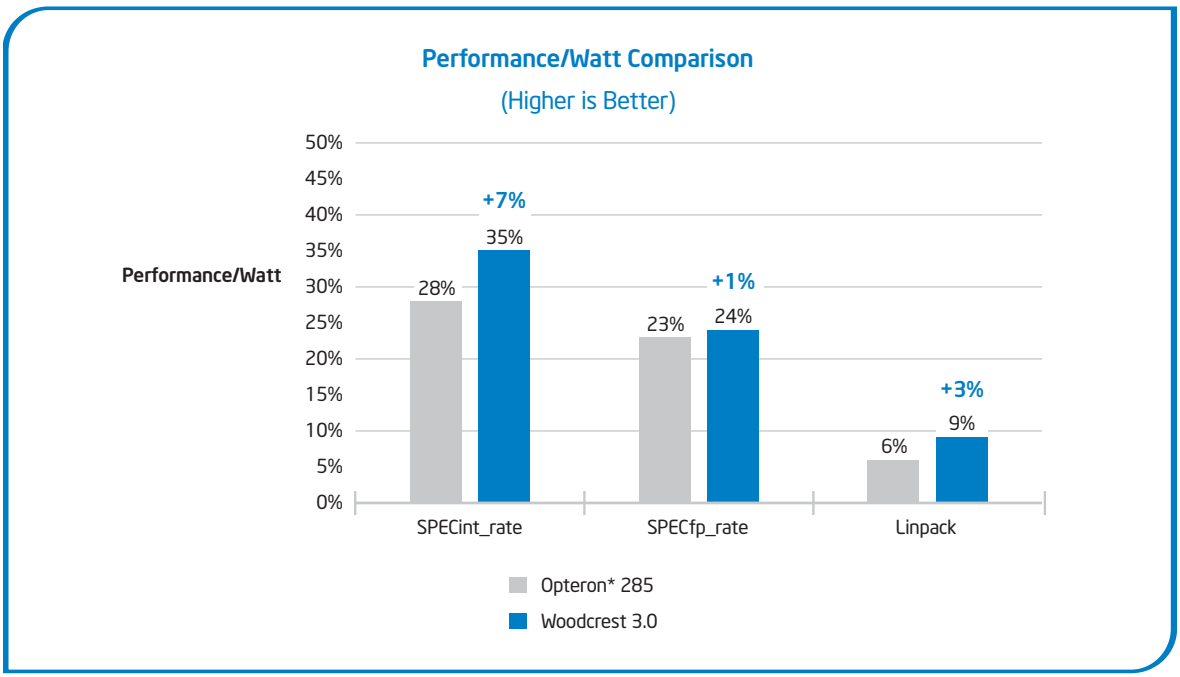

[4] Source: "Intel Woodcrest, AMD's Opteron and Sun's UltraSparc T1: Server CPU Shoot-out," by Johan De Gelas, AnandTech, June 7, 2006: www.anandtech.com/IT/showdoc.aspx?i=2772&p=11

[5] Based on Intel internal measurements made on May 3, 2006, comparing performance and power consumption with previous 64-bit Intel Xeon Processor based platforms. 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.

[6] For specific benchmark results, visit the Intel Web site at: www.intel.com/performance/server/xeon/index.htm

To help IT organizations balance power versus performance in diverse HPC 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 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.

Key advances include:

- **More instructions per clock-cycle**—Each core can process up to four simultaneous instructions, versus only 3 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.

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

Dual-Core Intel® Itanium® 2 Processor

Servers based on the upcoming Dual-Core Intel Itanium 2 processor offer comparable benefits for HPC applications that scale most effectively on large, shared-memory servers (Figure 2). They deliver up to 2 times the performance and 2.5 times the energy-efficiency of previous systems. They are available in a wide range of symmetric multiprocessor (SMP) configurations, from 2-way servers and blades, to large systems with up to 2,048 processors.

With its small, high-performing, energy-efficient processing core, Intel Itanium 2 processor designs will scale effectively in next-generation, multi-core designs. The high level of per-core parallelism (up to 6

Figure 2. The Dual-Core Intel Itanium 2 processor delivers major performance and energy-efficiency benefits for HPC applications that scale most effectively on large, shared memory systems.

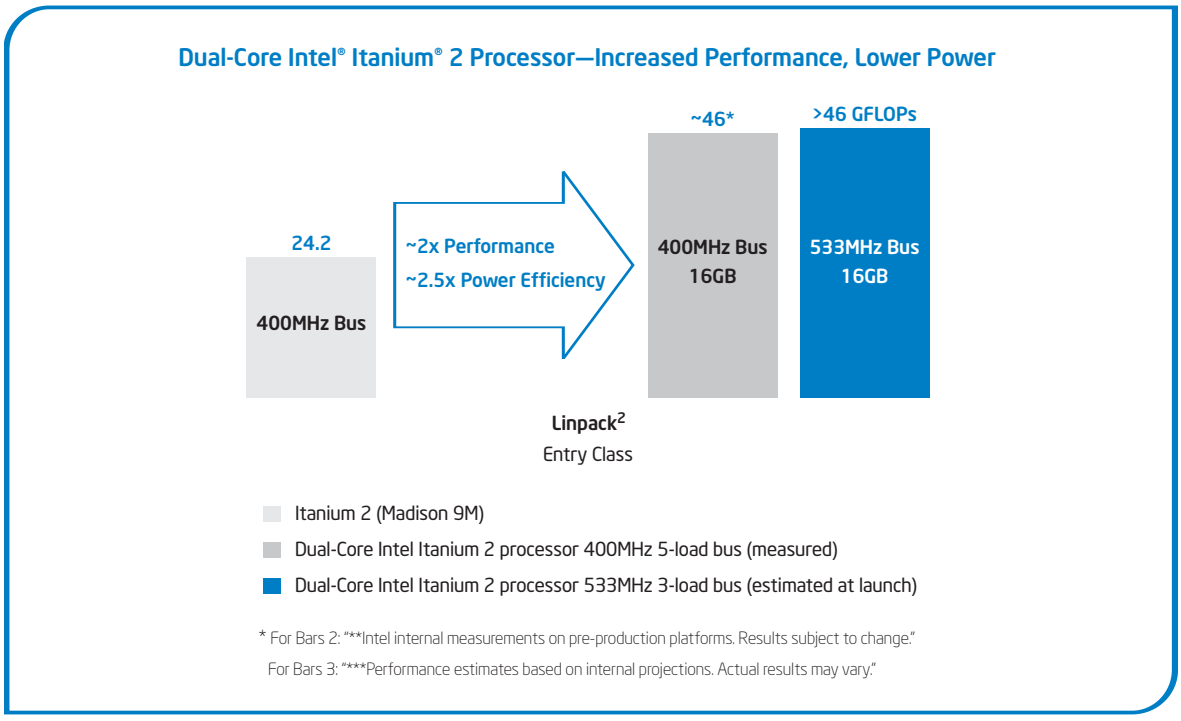

instructions per clock cycle with room for growth) will add to these advantages, especially in conjunction with ongoing compiler advances that continue to increase code-optimization for highly parallel throughput. These fundamental advantages will support rapid and continuing performance advances going forward, and help Itanium-based system vendors deliver leading performance per watt in comparison with competing RISC and mainframe systems.

Driving Parallelism with Multi-Core Designs

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 high-performance, energy-efficient foundation for current and future needs.

Rapid and Ongoing Gains in Performance/Watt

"We are targeting factor of ten improvements in both performance and energy-efficiency through the remainder of this decade, and are focused on delivering next-generation systems and software tools that will help organizations continuously scale the capability and value of their HPC environments."

– Pat Gelsinger, senior vice president and general manager, Digital Enterprise Group, Intel Corporation.

Intel has long-term technology roadmaps for both its Intel Xeon and Intel Itanium 2 processor families, and more than 10 processors with four or more cores are in development today. In all cases, there is an intense focus on delivering ongoing increases in performance and energy-efficiency through both per-core advances 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]

Optimizing System Performance per Watt

Though 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.

For example, compared to previous generations, new Dual-Core Intel Xeon processor-based servers provide:[8]

- **Up to 3x the system bandwidth**
- **Up to 4x the memory capacity**
- **Up to 3x the peak memory bandwidth.** Importantly, the memory subsystem operates at full speed even at full capacity, unlike competitive offerings, which require 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

[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

[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

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. Though 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 next-generation 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 are expected in the first half of 2007.)

The Critical Importance of Software

Multi-core processors deliver immediate performance and energy-efficiency benefits for optimized software, which include all major operating systems (OSs) and many existing HPC applications. However, many others are single-threaded, so they cannot take direct advantage of multiple cores. Optimizing these applications for multi-threaded throughput can accelerate performance, improve utilization and reduce the amount of energy required to accomplish the same work. 64-bit migration can deliver comparable benefits for many 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).

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). Low voltage Intel Itanium 2 processors extend this value, by enabling blade solutions that can handle the most data-intensive HPC applications.

- **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 leading OS and application developers 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://or1cedar.cps.intel.com/software college/) offer training, information and resources that can help HPC developers optimize their software more effectively and at less cost. Intel also offers extensive online resources for HPC developers at the Intel HPC Developer Center (www.intel.com/cd/ids/developer/asmo-na/eng/dc/hpc/index.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

The Future of Energy-Efficient Performance

Intel researchers are paving the way to an estimated 10x improvement in both performance and energy-efficiency over the next 10 years, through ongoing advances in silicon, processor and platform technologies.[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.

Figure 3. Intel is taking a comprehensive approach to solving today's power and cooling 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.

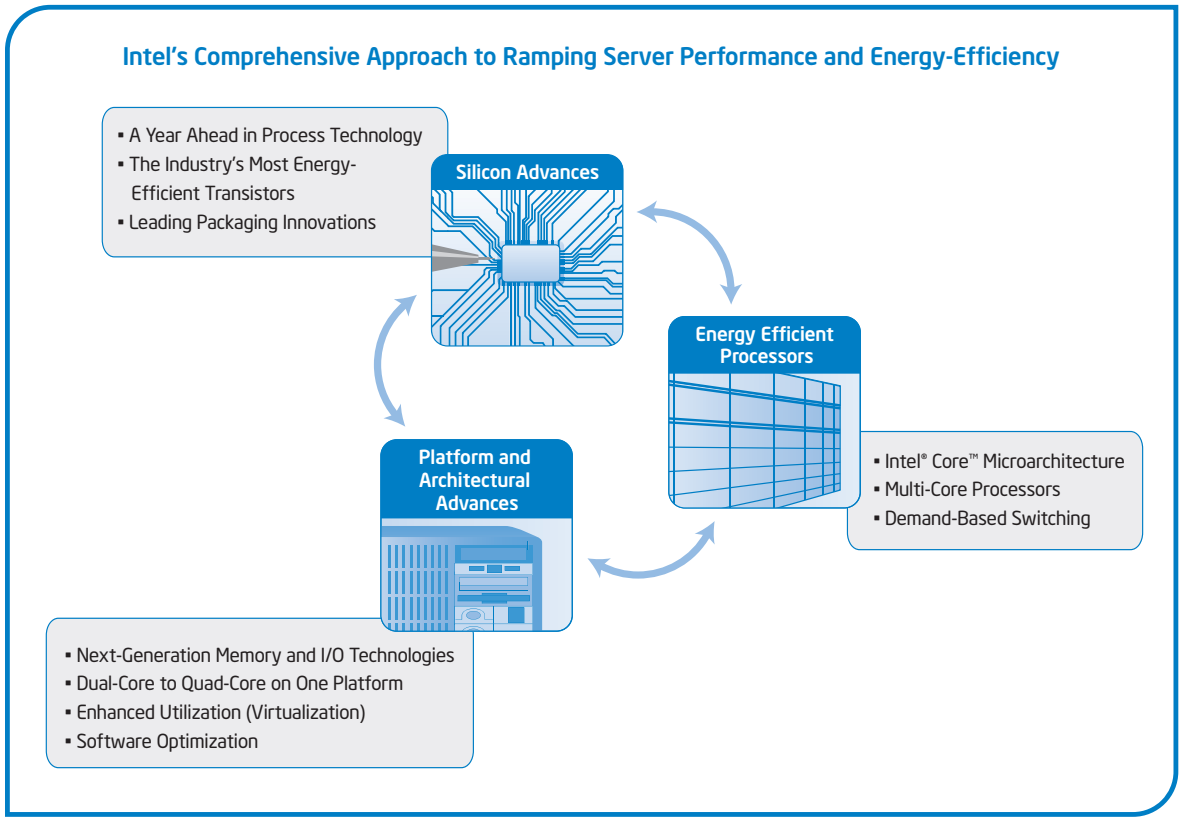

[13] *"Intel can talk not only about a road map of two cores, leading to four cores, leading to eight and so on, but also the manufacturing processes that will produce line widths of 65 nanometers, dropping to 45 nanometers, dropping to 32 nanometers, dropping to 22 nanometers. And, more to the point, 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/45nm_technology.htm

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. 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

Processor and Platform Innovation

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. 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 dual-core Intel processors, other platform components are contributing a higher percentage of the total power and cooling burden (Figure 4). Intel is working to increase energy-efficiency across the entire platform, and is leading the way toward system power consumption 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.

Figure 4. 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.

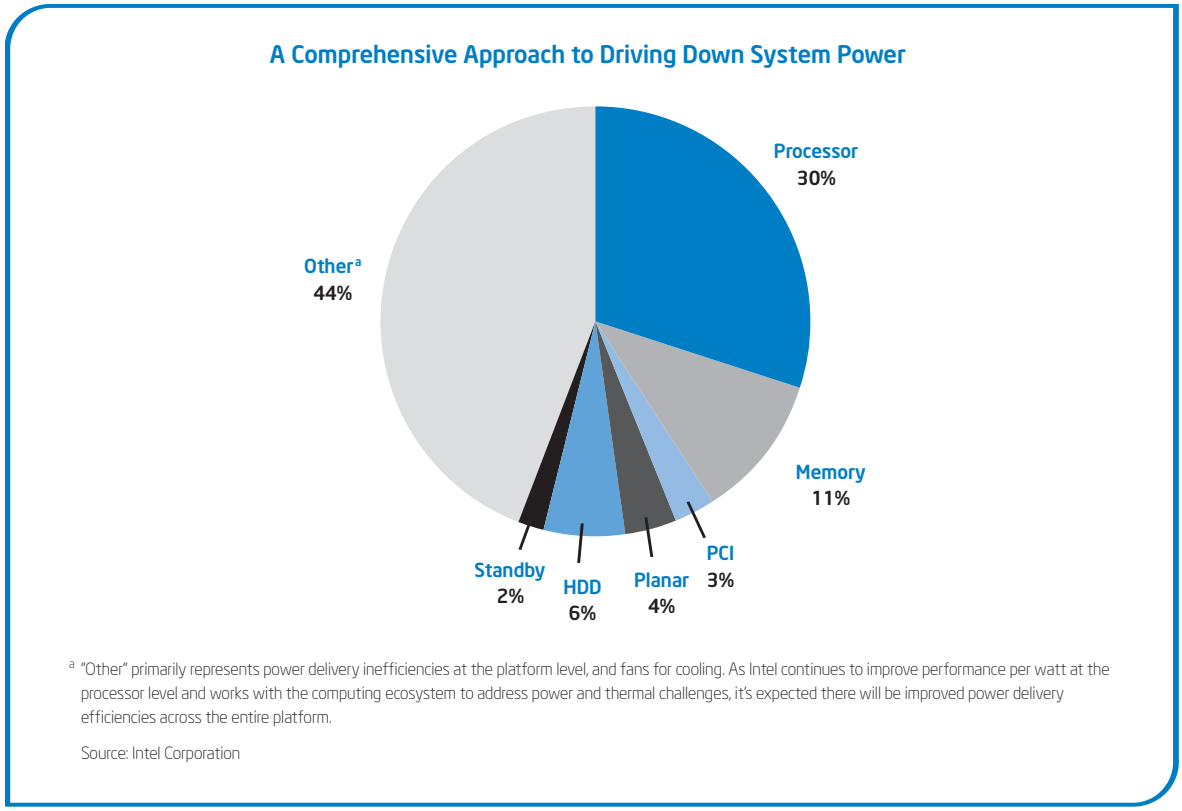

[a] "Other" primarily represents power delivery inefficiencies at the platform level, and fans for cooling. As Intel continues to improve performance per watt at the processor level and works with the computing ecosystem to address power and thermal challenges, it's expected there will be improved power delivery efficiencies across the entire platform.

Source: Intel Corporation

The Bigger Picture—Higher Performance at Lower Total Cost

"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]

Though increasingly important, optimizing energy-efficiency is only one factor in the larger goal of optimizing HPC capabilities continuously and cost-effectively. Performance, price/performance, interoperability, standards-compliance, software tools and support 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 platforms, and works with a broad range of hardware and software vendors, as well as open source developers, to drive comprehensive HPC advances that are closely aligned with emerging interconnect, clustering and grid technologies. The goal is to optimize total platform capabilities per watt and per dollar, and to deliver comprehensive tools and support. This helps to ensure that HPC solutions not only deliver better performance and energy-efficiency, but are also easier and less expensive to design, deploy, manage and scale. Based on these efforts, Intel processor-based servers have come to account for a full two-thirds of the world's largest supercomputers (see www.top500.org), and are widely used in HPC solutions of all sizes.

Just as importantly, 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 rapidly increasing performance, density and energy efficiency with every new generation—all on a consistent, standards-based architecture that helps them streamline operations and reduce TCO.

[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_gci1186068,00.html?track=NL-48&ad=545073&Offer=52HTMLA

Conclusion

New Dual-Core Intel Xeon processor-based servers provide breakthrough value in performance, price/performance and energy-efficiency. Upcoming Dual-Core Intel Itanium 2 processor-based systems offer comparable advantages for HPC applications that scale most effectively on large, shared-memory servers.

These new server platforms are the first step in a comprehensive drive to ramp server performance and density, while helping organizations drive down their power and cooling requirements. Through comprehensive platform innovation, software optimization and next-generation silicon process technologies Intel will continue to help organizations expand their HPC 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 solutions.

Additional Resources

- Accelerate your transition to dual-core Intel processors. Talk with your local Intel representative, and visit us on the Web, at: www.intel.com/go/xeon.
- Learn more about Intel products, technologies and resources for high performance computing, at: www.intel.com/go/hpc.

Appendix A: Performance and Energy-Efficiency Tests and Comparisons

New Dual-Core 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

Appendix B: Additional Resources

Intel Power & Cooling Online 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 IT costs and options for many years. Effective solutions must take into account not only changing 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 professional services organization within Intel, offers expert assistance based on formal, best-known methods for optimizing the use of current facilities, projecting future requirements, and developing long-term strategic plans that balance opportunity against TCO. Intel Solution Services has worked successfully with many clients to develop detailed technical and financial analyses that drive better and faster decisions on crucial data center issues. For more information, visit: www.intel.com/go/intelsolutionservices

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/
IBM Intel-Based Solutions: www-03.ibm.com/systems/x/
Microsoft: www.microsoft.com/management/default.mspx

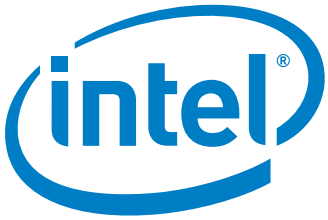

www.intel.com

0606/JRR/IMP/XX 313637-001US