Five Ways to Maximize Web Tier Efficiency

White Paper
Intel® Xeon® Processors
Digital Enterprise

Introduction

Explosive Internet growth and richer applications are putting constant pressure on Web-based companies to find ways to scale their infrastructure at the lowest cost. Nowhere is this pressure more evident than in the Web tier, the front line servers that respond to client requests.

In an Internet data center, the Web tier consists of the servers and applications that interface most directly with the client in the delivery of Internet services. These front-line servers typically outnumber other data center servers by as much as a factor of five. Consequently, it's also the area in which great cost savings can be achieved through energy-saving improvements. Depending on the size of the operation, the Web tier in a data center can consist of hundreds to thousands of front-end Web servers. Consequently, saving a handful of watts per server can add up to tens of thousands of dollars over the useful life of the infrastructure. Consolidating servers with more capable, efficient servers can have an even larger impact.

To help you realize these savings, this white paper provides an overview of five strategies for maximizing Web tier efficiency through more efficient server technologies.

Table of Contents

Strategy #1: Evaluate Servers Based On Performance and Efficiency In The Real World.

A long-time challenge in the selection of energy-efficient servers has been the lack of an appropriate metric for comparing performance and efficiency on the basis of real world loads. Most efficiency metrics compare performance-per-watt at just a single point along the load line—such as at idle or at 100-percent CPU utilization. Neither of these points are an accurate measure of real world usage. The amount of work a server does typically varies greatly depending on its function and usage through the day.

Intel has worked with industry-standards bodies such as the Standard Performance Evaluation Corporation (SPEC) and The Green Grid to develop comprehensive and unbiased performance efficiency benchmarks. One such benchmark is SPECpower.* Using this benchmark, one can measure the performance-to-power ratio from idle to 100 percent of the target load and all points in between. SPECpower is the first and so far only industry standard server energy efficiency benchmark.

SPECpower is expressed as a performance-to-power ratio, so the higher the number, the better. The benchmark measures how a server executes a standard Java* workload at CPU load levels in 10-percent increments from 0 percent to 100 percent. The results are published in text and graphical form so users can see a chart of how each server performs. Instead of trying to judge power consumption by the wattage listed on the back of a server, SPECpower gives a real idea of the work that can be done at specific energy consumption level.

Using this benchmark, buyers for Web tier applications can compare servers from different manufacturers that use the same processors and determine which provide the optimal performance to power.[1] According to Andrew Fanara, the director of the U.S. Environmental Protection Agency's (EPA) ENERGY STAR* Product Specifications Development team: "SPEC has taken a critical first step to give server vendors and their customers a standardized benchmark tool that elevates power efficiency in the performance evaluation process."[2]

SPECpower has been designed to be easy to set up and run. Running the benchmark requires a server (the test system), a suitably accurate power analyzer, thermal sensor, and a control system (such as a notebook). Find out more about the SPECpower benchmark at ***http://www.spec.org/power_ssj2008/***.

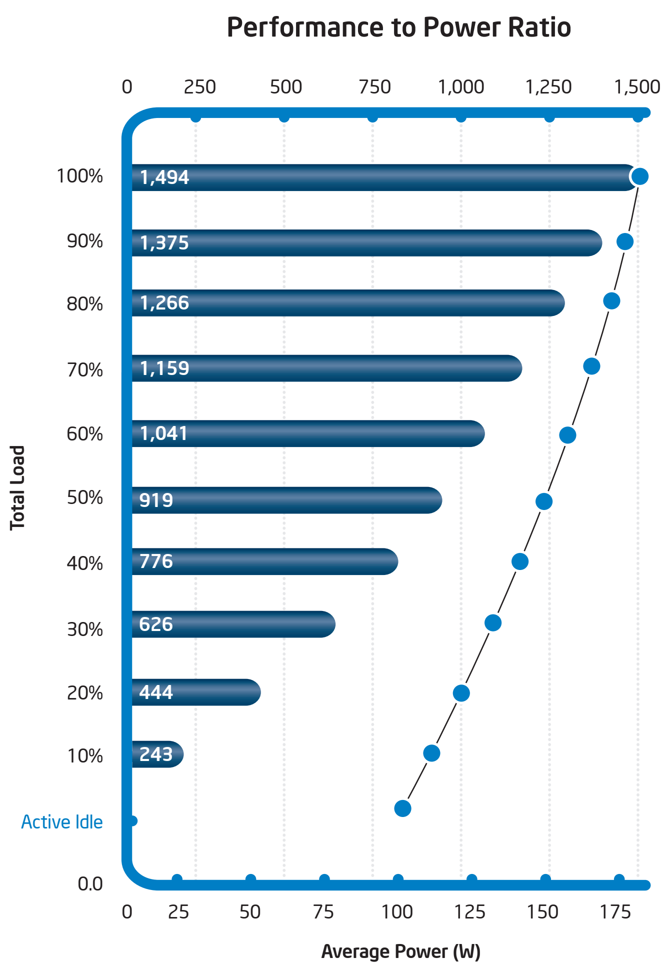

Figure 1: Efficiency should be measured by both performance and real world power loads, as shown in this sample SPECpower chart.

Strategy #2: Select Best-In-Class Server Components.

Virtually every component in a server draws a certain amount of power, some more than others. (See Figure 2.) One obvious way to reduce power consumption is to purchase servers with more energy-efficient components. Take power supplies, for instance. They are generally about 70–75 percent efficient and can account for as much as 21 percent of a server's power consumption. At 75 percent efficiency, they require 100 W of AC input to produce 75 W of DC output and dissipate the remaining 25 W in heat. Higher-quality power supplies are available that can be over 90 percent efficient. These use less direct power and also require less power to cool.

To maximize energy efficiency, ask your server vendor for the best ways to reduce power consumption. Use the following component selection tips as a starting point.

1. Specify fewer, but larger capacity memory DIMMs. In general, the more memory DIMMs installed, the more power they will consume. For example, a server with sixteen 1 GB FB-DIMMs could consume up to ~160 W under load versus ~44 W for a configuration using four 4 GB FB-DIMMs. In this case, choosing fewer, but larger capacity DIMMs could save up to 116 W per server.

2. Compare the power consumption of various brands of DDR2 or FB-DIMM memory components. DIMMs can vary by as much as 5 W for the same memory capacity depending upon vendor/model.
3. Look for Climate Savers Computing Initiative-qualified systems with 85 percent-plus efficiency PSUs. Moving from a 75 percent efficient power supply to an 85 percent efficient power supply could save up to 25 W on a 250 W server.
4. Specify server boards with feature sets optimized for the Web tier. For example, avoid boards with features such as multiple PCI Express* slots or I/O connectors that are often unnecessary in Internet scale deployments.
5. Consider solid state drives (SSDs) in place of traditional hard drives. A typical SAS (serial-attached SCSI) hard drive consumes ~12 W under load versus ~2 W for SSDs. Replacing two hard drives in each server with SSDs could save 20 W.
6. Ensure that the server vendor enables all the power management options in BIOS, particularly Enhanced Intel SpeedStep® technology, enhanced HALT state (C1E), and Closed Loop ThermalThrottling (CLTT).

Component Level Power Consumption in a Typical Server

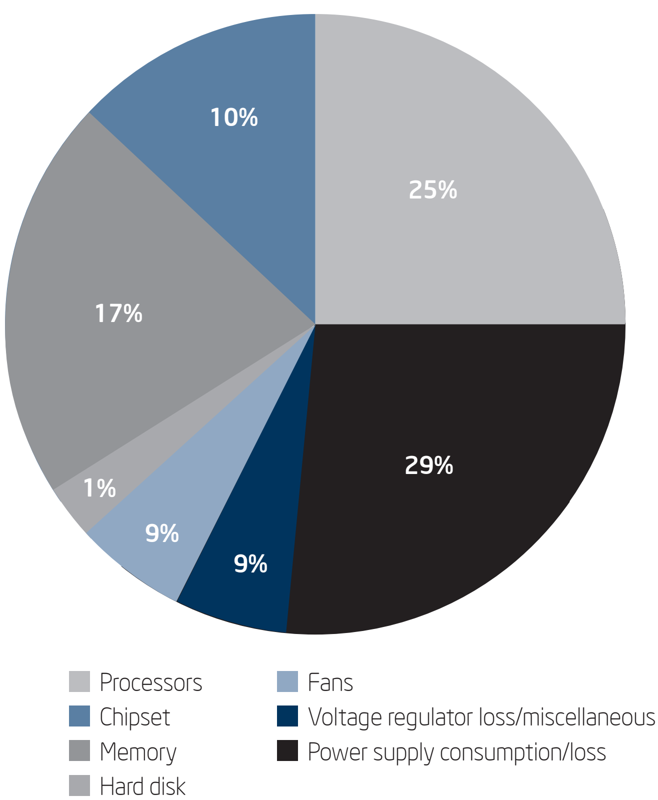

Figure 2: Component level power consumption in a typical server, using two 50 W low voltage processors.

Source: Intel internal estimates as of October 2008 using two 50 W processors, eight DIMMs, one hard drive, and one power supply.

How much of a difference can following tips like these make? A lot. To demonstrate, Intel assembled an Eco-Rack Proof of Concept using many of these energy-saving measures. Developed at Intel facilities in both Hillsboro, Oregon and DuPont, Washington, and inspired by discussions with the EPA and Lawrence Berkeley National Laboratory (LBNL), the Eco-Rack consisted of a test setup of 30 rack-mounted servers. The demonstration achieved 23-30 percent total power savings, depending on the workload, as measured by running compute-intensive enterprise applications such as SPECjbb* (a warehousing application simulating many transactions at the same time).[3] To learn more about the Eco-Rack and the results of each individual power management option and energy-saving measure, download the white paper "Intel Eco-Rack 1.5" at:

http://download.intel.com/technology/eep/ecorackwp.pdf.

Strategy #3: Take Advantage of Server Instrumentation for Power Management.

Another recent technology for improving power utilization is server instrumentation for power management. This new technology is available in selected server products today, but is expected to become more mainstream in the next year or two. Dynamic power management control is now available that enables a server to report power usage data and have platform power dynamically adjusted. Such information and control enable data center managers to: 1) match power allocation to actual server power consumption; and 2) more accurately populate a rack for maximum value of the rack space.

The Intel® Dynamic Power Node Manager, for instance, is an out-of-band (OOB) power management policy engine embedded in Intel server chipsets starting with next generation Intel® server platforms code-named Nehalem-EP. Node Manager works with BIOS and operating system power management (OSPM) to dynamically adjust platform power to achieve maximum performance/power at node (server) level. Node Manager has the following features:

- Dynamic power monitoring – measures real-time power consumption data (point in time, or average over an interval), and reports through the Intelligent Platform Management Interface (IPMI).
- Platform power capping – sets platform power to a targeted power budget while maintaining maximum performance for the given power level.
- Power threshold alerting – monitors platform power against targeted power budget. When the target power budget cannot be maintained, Node Manager sends out alerts to the management console.

Using Node Manager, data center operators can look at a typical workload and assign a budget, say limiting the system to 340 W. If servers demand more power, performance of secondary systems slows down to stay under the preset power limit. Alerts can be set to automatically notify staff when certain power limits are challenged.

Node Manager helps significantly improve rack density, increase performance density, implement load balancing improvements between more efficient and less efficient servers, and improve overall availability and total cost of ownership. (See sidebar.)

Saving Power with Intel Dynamic Power Node Manager Technology

Intel's Digital Enterprise Group worked with Baidu.com—a leading search company in China—to conduct a proof of concept (POC) project using the Intel Dynamic Power Node Manager Technology (Node Manager) to dynamically optimize server performance and power consumption to maximize the server density of a rack. The POC engineers used Node Manager to identify optimal control points, which became the basis to set power optimization policies at the node level. A management console—Intel® Datacenter Manager (Datacenter Manager)—was used to manage servers at rack level to coordinate power and performance optimization between servers to ensure maximum server density and perform yield for given power envelope for the rack. The POC was conducted in early 2008 at a Intel-Baidu joint lab under a simulated Baidu production environment. The POC results[4] showed that:

- At a single node level, **up to 40 W savings** per system without performance impact when an optimal power management policy is applied.
- At rack level, **up to 20 percent additional capacity** increase within the same rack-level power envelope when aggregated optimal power management policy is applied.
- Compared with today's data center operation at Baidu, by using Node Manager there could be a **rack density increase of up to 20 – 40 percent**.

Strategy #4: Optimize Application Software.

Optimizing code and using the latest compilers can be a very effective way to improve application performance and reduce the size of an Internet data center's server infrastructure. Often—when a Web company is starting out—the highest priority is to launch with stable code, and performance optimizations are a lower priority. Yet as the company grows, the complexity of the code grows as well as the size of the server infrastructure. This presents an opportune time to analyze the code for efficiencies. Some key areas to consider include:

- Optimizing on the latest Intel®-based compiler.
- Increasing threading support for multi-core processors.
- Identifying bottlenecks in the code with a performance analysis tool.
- Leveraging large memory configurations with 64 bit code.

In some cases, streamlining code can reduce the number of servers required to support Internet demands. This drives down both capital and operations expenses in the data center.

Intel offers software tools to enable companies to analyze and improve their software. Intel® VTune™ Performance Analyzer, for example, enables developers to identify performance bottlenecks without recompilation and with very low overhead. Intel® Threading Analysis Tools can help create correct code by vastly simplify the debugging and performance optimization of threaded code. More information about Intel's optimization tools is available at ***http://softwarecommunity.intel.com***.

The gain can be substantial. Sapan Panigrahi, Senior Manager of Yahoo Search Performance Engineering, notes: *"We have been using VTune analyzer for last two years. VTune has proved to be an invaluable tool in optimizing our business-critical applications by isolating performance bottlenecks and providing insight into techniques of removing these bottlenecks. It has helped us realize significant return in our hardware investment."*

Strategy #5: Select A Processor That Delivers The Best Energy-Efficient Performance.

Inside each Web server, of course, is one or more processors. It pays to give particular attention to which processor is selected because processors can consume a sizeable percentage of a server's overall power. Paradoxically, processors are not the cause of the skyrocketing energy costs in the Web tier. In fact, processor advancements continue to be one of the greatest sources of energy savings within the data center. The problem lies with the sheer number of servers in the Web tier and often their age. To reap the performance and power benefits of the latest server processors, an organization must consider replacing older servers on a regular basis. For example, Intel IT has standardized a server refresh cycle of four years to take advantage of the huge performance improvements and power reductions offered by latest energy-efficient processors.

How great are these benefits? Consider Moore's Law, the well-known prediction made by Intel co-founder Gordon Moore in 1965 that the number of transistors in a chip would double about every two years. Through a long history of Intel innovations, this observation has continually resulted in dramatic improvements in transistor density resulting in both greater core integer performance and power reduction over time. In fact, over the past 30 years, the relative energy-per-transistor ratio has improved approximately one million times.

In the age of mainstream multi-core processing, the gains in performance for unit of energy have been especially dramatic. As Intel® Xeon® processors have improved from single-core to multi-core processing, performance-per-watt has increased by as much as eight-fold within an essentially flat system power envelope. In other words, in approximately three years, it's become possible to achieve as much as eight times greater performance for the same amount of power.

Let's put this in server terms. In just three years, the processing power of six racks of servers now fits in a single rack and consumes one eighth the energy (see figure 5). Obviously, switching to servers using more efficient processors enables the increase of processing capacity while decreasing energy consumption. More requests can be handled for less energy.

Power Reduction Over Time*

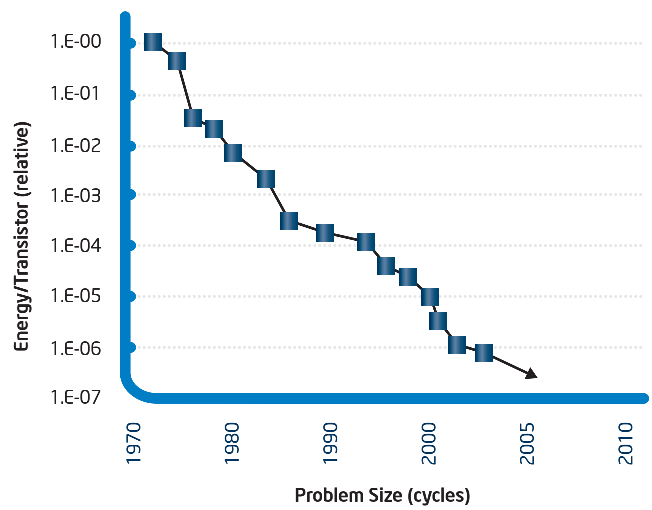

Core Integer Performance Over Time*

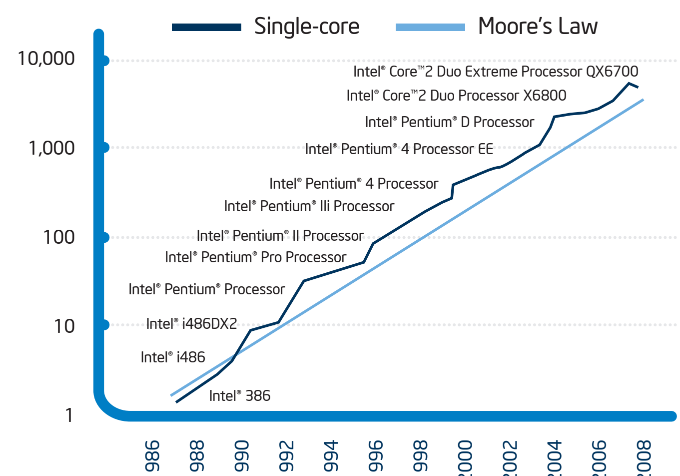

Figure 3: Benefits of Moore's Law.

The Benefits Of The Latest Gains In Processor Efficiency

A good example of the latest energy-efficient processors is the 50 W 45nm Intel® Xeon® Processor 5400 series. Featuring enhanced Intel® Core™ microarchitecture, these 45nm high-k quad-core processors include a large 12 MB on-die cache for fast response and deliver a host of energy-saving benefits for the Web tier.

- Up to 25 percent higher performance[7] than prior generation low voltage Intel Xeon processors.
- A low 16 W idle power to reduce average system power and operating costs.
- 35 percent less processor power consumption compared to 80 W TDP mainstream processors means more servers per rack.

Simply specifying that new servers be based on these and other Intel energy-efficient processors goes a long way in increasing processing power and making better use of available power.

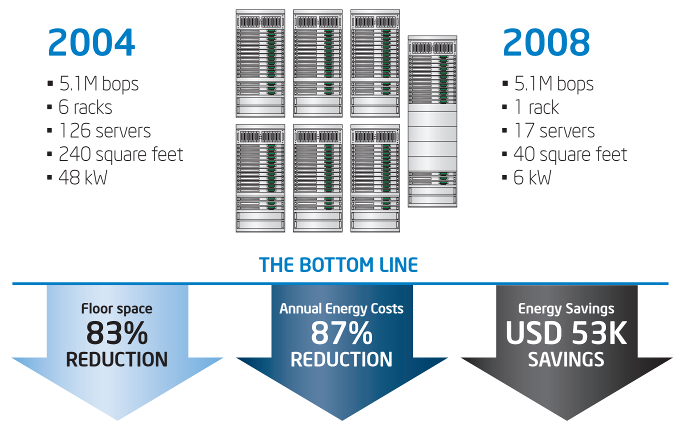

Figure 5: Processor innovation enables deployment of fewer racks to do an equal amount of work.

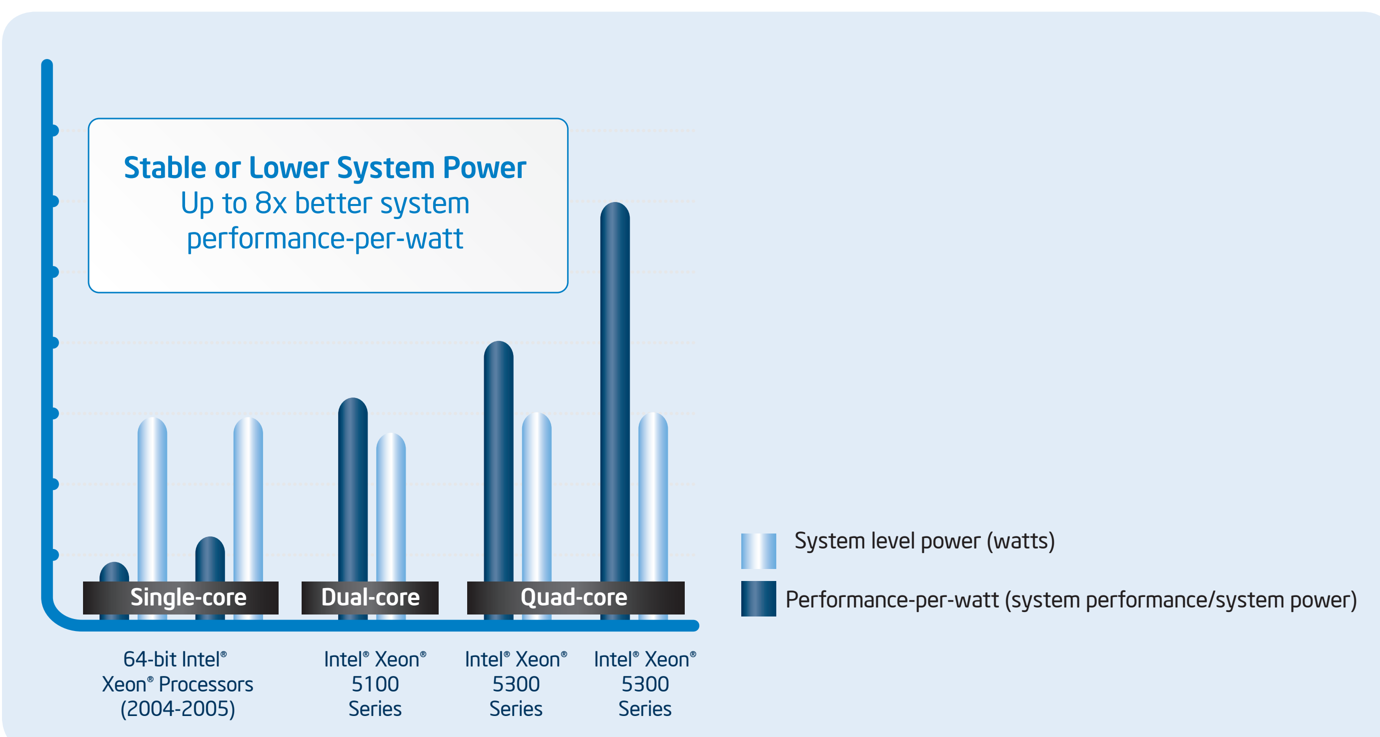

Figure 4: Intel processor innovations have dramatically boosted performance-per-watt over the last three years.

Conclusion

Few businesses face greater server scaling challenges than companies offering large-scale Web services. The good news is that through industry-wide efforts led by companies like Intel, industry organizations like The Green Grid, and government agencies like the EPA, more ways to improve and measure server performance per energy unit under real world loads are now available. By taking advantage of these technologies and measurement tools, Internet businesses can significantly improve performance-per-watt and engineer real savings throughout their Web tier or increase capacity within power- and space-constrained environments. What's more, these businesses can take advantage of highly efficient Intel Xeon processors today for large scale data center deployments—processors which have garnered the highest scores in the industry standard SPECpower benchmark.[8] Armed with industry standard measurement tools, optimization techniques, and Intel's industry leading technology, companies have some outstanding resources for achieving greater efficiency in their large scale data center deployments.

For more information on performance tests and on the performance of Intel products, visit ***www.intel.com/performance/resources/limits.htm*** or call (U.S.) 800-628-8686 or 916-356-3104.

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 effect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit www.intel.com/performance/resources/limits.htm or call (U.S.) 800-628-8686 or 916-356-3104.

1. See http://www.spec.org/power_ssj2008/results/power_ssj2008.html for a list of currently tested servers and the results.
2. SPEC "benchmark measures servers' performance-to-power ratio," SearchDataCenter.com, December 12, 2007.
3. For a full description of the equipment tested and compared, see the white paper "Intel Eco-Rack 1.5" at http://download.intel.com/technology/eep/ecorackwp.pdf.
4. For a full description of this proof of concept and how its results were achieved, see the white paper "Dynamic Power Optimization for Higher Server Density Racks – A Baidu Case Study with Intel® Dynamic Power Technology" at http://communities.intel.com/servlet/JiveServlet/previewBody/1492-102-1-1723/Node%20Manager%20Baidu%20POC%20WhitePaper%20-%20External.pdf;jsessionid=C452E6DC3A0B64D3F4F2BD2456DAD1FE
5. Performance as measured by results running SPECjbb2005.* For details on systems tested and other test specifics, see: http://www.intel.com/performance/server/xeon/ppw.htm
6. Source: Intel January 16, 2008. Performance comparison using SPECjbb2005 bops (business operations per second) between (2004) 2 socket Intel® Xeon® processor (3.6GHz) and (2008) 2 socket Intel® Xeon® processor E5450 (3.0GHz, 80 W, 1333 MHz) measured 8/22/07. 1 Floor space based foot on sq ft reduction. 2 Energy costs based on total solution rack power. 3 Energy savings based on an electric rate of $0.10/ kWh assuming 33% average server load. 4 ROI is calculated based on cost of new servers ($6,264 per server pricing based on HP DL 380G5 (32 GB RAM) as of Nov 17, 2007, source www.hp.com) divided by energy savings per year. Actual performance results and savings may vary depending on configuration.
7. 25% claim based on performance comparison on SPECpower*_ssj2008 benchmark. Comparison details: Baseline: SuperMicro X7DBE*+ Server platform with 2 x Intel® Xeon® processor L5335 2.0GHz, 8 M L2 Cache, 1333 MHz FSB, Intel® 5000P chipset, 4x2GB FB-DIMM 667 MHz memory, Microsoft Windows Server 2003 Enterprise x64 Edition* SP1, BEA JRockit* JVM. New platform measurement: Same configuration as the baseline except with 2 x Intel® Xeon® processor L5420 2.50 GHz, 12 M L2 Cache. Intel internal measurement. March 2008.
8. Performance and energy-efficiency leadership as measured by the highest score on the SPECpower_ssj2008 benchmark for a 2 socket server, demonstrated by an NEC ECOCENTER system with the Intel Xeon L5420 processor. Source: Public SPECpower results from http://www.spec.org/power_ssj2008/results/power_ssj2008.html as of September 2, 2008.

Printed in USA 1008/VP/HBD/PDF