Benefits of Dual-Core Computing on IBM System x Servers Using Intel Processors

By Mark T. Chapman, Cristian Rojas, and Maurice Bland
IBM Systems and Technology Group

Executive Overview

This paper examines the real-world advantages of dual-core processors, using the new IBM® System x™ servers as real-world examples. We discuss the benefits to industry-standard server platform users, and demonstrate scalable performance and superior performance-per-watt for high-performance HPC and business computing applications.

Following the ongoing transition to 64-bit computing, the next step in advanced processor technology was the introduction of dual-core processors, followed by the new generation of lower-power dual-core processors. In the past, typical processor chips contained one central processing unit, or CPU core (the "brains" of the processor), surrounded by supporting circuitry, such as on-chip L2 cache. A "2-way" server in this case would have two processor sockets, each containing one single-core chip. By contrast, dual-core processors contain two complete processor cores within one chip, along with L1 and L2 cache and other supporting circuitry. This offers the potential for greater performance and reduced latency, along with lower power draw and heat output, than two physical processors[1] would provide.

This paper examines the potential advantages of Intel® *dual*-core processors over Intel *single*-core processors in IBM System x platforms. The two metrics we will use are *overall performance gains* and *performance-per-watt*. Performance-per-watt measures how much "work" can be done for every watt of power used. This is an extremely important metric because as businesses grow, so does their need for a more efficient infrastructure.

Our results show *performance gains* of **30-110%**, as well as a **90%-175%** increase in *performance-per-watt* when comparing an IBM System x3650 (dual-core) server to the predecessor IBM eServer™ xSeries® 346 (single-core) server[2].

---

[1] Physical processors of the same type and clock rate.

[2] Although the benchmark results shown are specific to the x3650 server, the same processor, memory, and I/O technologies are used in other x3XXX series servers. Therefore, similar relative performance enhancements can be expected when compared with predecessor IBM products.

Introduction

Why Dual Core?

Since the invention of the microprocessor approximately 35 years ago, computer engineers and architects have found innovative ways to use the transistors at their disposal to create faster and cheaper devices. The first commercially available single-chip microprocessor, the Intel 4004, had a 4-bit central processing unit (CPU), operated at less than 1 MHz, and was built on a 10-micron process with about 2,300 transistors. Today, a 64-bit microprocessor operates above 2GHz, is built on sub-100 nanometer processor technology, and uses several hundred million transistors.

The extremely fast rate and pace in the progress of processor technology over the past four decades was first observed and predicted by Gordon Moore, an Intel co-founder, and has commonly been referred to as Moore's Law. Specifically, Moore's Law states that the advancement in the manufacture of electronic integrated circuit packaging technology makes it possible to double the number of transistors per unit area every 18 to 24 months. The practical effects of Moore's Law have been to give technology users ever-increasing compute capabilities while *decreasing* the cost. Over the years, Moore's Law has made it easy for chip manufacturers to increase processor performance. Along with the advancements in each generation of manufacturing process technology, transistors have become smaller and more densely packaged, thus reducing capacitance and wiring distance between circuit interconnects. This has allowed processor clock frequencies to soar.

Today, anyone who follows the IT industry knows it is increasingly difficult to scale up the clock rate of processors. The benefits of Moore's Law come with a "dark side": as transistor densities increase and processor frequencies climb, the amount of wasted power and heat produced by the processor also climbs. The corresponding increase in processor power density has made it increasingly difficult to remove the generated heat using traditional low-cost system cooling techniques. This has become a major factor in limiting how fast a processor can go. Power density is typically represented as *power per unit area* or in *watts/cm*^{2}, and if we were to extrapolate processor power density over time at the same rate and pace experienced by the industry over the past 20 years, the ever-shrinking area of the processor would eventually produce more heat than the surface of the sun! The higher the clock frequency, the more wasted heat it generates. The hotter the chip, the hotter the system runs and the more challenging it becomes to keep everything cool.

While smaller transistors can operate at lower voltages, which results in less power consumption and less heat produced, this decrease in transistor size will also *increase* device leakage current. Current leakage is responsible for as much as 40% of total device power consumption, driven by gate leakage and off-state leakage current—the continued flow of electrons through the transistor dielectric "wall" even when the transistor is switched "off." Leveraging transistor voltage threshold (Vt) is one design technique used for reducing leakage. **Low Vt** transistors are fast but have high leakage, while **high Vt** transistors have low leakage but are slower. Balanced performance at low power can be achieved by using low Vt transistors in a small number of performance-critical locations, while primarily using medium or high Vt transistors to minimize leakage.

In the end, the number of transistors per unit area has increased to allow more sophisticated processor implementations, but the operating frequency has had to decrease in order to be able to keep the processor power envelope at a manageable level. (This is why there are no 4GHz processors available from Intel and other vendors.)

What is needed is other manufacturing and design approaches that focus on *processor performance per watt*. Microprocessor designers are now looking for ways to constrain power at all levels. Processor architects, circuit designers, and process engineers are focusing on reducing leakage and making power-efficient solutions. This is being addressed through manufacturing process technology, in creative circuit design and microarchitecture, in power management features enabled in the operating system, and in greater efficiency in system power distribution and delivery.

The historic breakneck rate of processor frequency scaling that technology users have come to expect as a side effect of Moore's Law cannot be maintained. Blindly driving GHz as the key performance metric has now challenged the limits of physics, due to the exponential rise in power consumption and power density. By shifting from single-core processor technology to dual-core, manufacturers have been able to step down the clock rate considerably, while still increasing performance. Going forward, frequency increases of about 10% per year are all that should be expected, as a result of future manufacturing process technology advancements.

Moore's Law has not been broken; processor chip manufacturing process technologies are expected to continue enabling higher transistor densities. Instead of ramping up clock frequencies to drive higher performance, processor vendors will continue to find ways to increase the number of processor cores on the die. Dual-core processors are here today, with quad-core processors expected to ship by 2007. Designs for eight-core and higher processors are already planned. As transistor sizes shrink and the number of cores per chip increases, so will the size and sophistication of the on-board cache memory available to those cores. This can yield additional performance benefits. These advancements are discussed in more detail in the following sections.

Dual-Core Benefits

Two major advantages gained by switching from single-core to dual-core processors are greater overall performance and improved performance-per-watt efficiency.

In the past, the simplest alternative was to increase the size of the internal L2 cache and perhaps add an external L3 cache. This allowed more of the program code to reside near the processor, reducing the need for comparatively slow accesses of main memory (RAM). Processor vendors have been doing this for years. Another method has been to increase the number of processor registers, so that more processor instructions can be available simultaneously. However, these techniques alone will not allow for processor throughput to continue accelerating at the same rate it has been.

In a system designed with multiple processor sockets, a second processor can be added. This often increases performance—but not by 100%—due to resource contention and latency issues. The more processor sockets that are added, from two to four, from four to eight, and so on, the more intractable the contention and latency issues become.

The most recent performance-enhancing technique is to fit a second processor core and associated L2 cache inside one physical chip, creating a "dual-core" processor. Placing two cores close together in a single socket eliminates much of the interprocessor latency found in a two-socket SMP system. In essence, a dual-core processor is a "2-way SMP" system on a single chip. This design allows a *dual*-core processor chip, running at a lower clock speed, to outperform (for many applications) a similar *single*-core chip running at a somewhat higher clock rate.

There is more to processor performance than merely clock rate, bus speed and cache size. Another important aspect is how effectively the processor handles program threads (and how "threaded" the software stack is). Take, for example, a *single*-core, dual-threaded processor. Each thread is assigned its own set of registers. This makes the processor appear as two (logical) processors. Theoretically, it can process two threads indefinitely. In reality, this rarely happens because those logical processors are still part of one physical processor, requiring the threads to time-share common resources such as integer units, floating-point units and cache. This means that there can be contention between the threads for the same processor resources. As a result, the effective throughput is somewhat less than the theoretical, with one thread waiting for the other to release a shared resource. In addition, if the software stack is largely single-threaded, the second logical processor may be idle much of the time, rendering its value moot.

By contrast, with a *dual*-core dual-threaded processor, two physical processors reside inside one chip. Because each core has its own cache, registers and other resources, there is less resource contention than you might see with a simple dual-threaded, single-core processor. Two separate single- or multi-threaded programs can be running simultaneously, for up to twice the throughput of a same-speed, single-core processor.

For most server applications, despite running at a lower clock rate, a 2.33GHz *dual*-core Intel Xeon® processor can offer significantly greater total throughput than a 3.6GHz *single*-core processor.

Another advantage of dual-core processors is the power and thermal reduction. One *dual*-core chip offers approximately the same performance as two *single*-core chips of the same clock rate, while using half the power and producing half the wasted heat. This can save you a significant amount of money over the long term. In today's datacenter environment, a server's performance per watt of power is becoming an ever-greater concern for IT managers. Dual-core processors can improve this situation considerably.

By stepping the clock rate back a few notches in dual-core processors, engineers have managed to increase performance while pushing back the date when higher clock frequencies become an insurmountable roadblock with current technology. With IBM—as far back as 2002—and more recently Sun, AMD, and Intel, all adopting dual-core processor designs, we are unlikely to see any new single-core server processors after 2006. It is not a question of *whether* you will migrate to dual-core processors, merely how soon.

So what of software licensing issues? Will the move to dual-core processors mean higher software costs? Probably not. Microsoft® and many other software vendors have announced that they plan to license software according to the number of *physical processor sockets*, rather than the number of cores; therefore, the transition from single-core to dual-core processors should have little or no negative effect on software licensing fees for most customers. In fact, by doing the work of two single-core servers today, dual-core servers can help delay the day when a second physical server will be needed.

Note: As with most performance enhancements, this one comes with the caveat that some customers and some tasks will benefit from it more than others. Just as adding memory only helps if your applications can use it, dual-core processors will help those users with multithreaded, compute-intensive applications more than those with single-threaded or I/O-intensive applications.

Will You Benefit?

A number of categories of applications and middleware can benefit from dual-core processors. To help you determine whether these new processors will help you, *Table 1* summarizes the categories and representative applications that could see significant benefit, moderate benefit, or little benefit from a switch from single-core to dual-core processors. (The categories are listed in approximate order of greater-to-lesser benefit within each column. The applications in each category—such as HPC servers—are listed in alphabetical order.)

Significant Benefit	Moderate Benefit	Little Benefit
HPC Servers • Automotive • Aeronautical • EDA • Geophysical • Life sciences	**Web Servers** • Apache • Microsoft Internet Information Server	**Terminal Servers** • Citrix • Microsoft Terminal Server (32-bit version)

Significant Benefit	Moderate Benefit	Little Benefit
DCC Servers • Pixar RenderMan	**File Servers** • CFIS • NFS • Samba	**Legacy non-threaded server applications**
Data Mining • MicroStrategy • SAS	**E-Mail Servers** • Lotus Notes • Microsoft Exchange • Sendmail	
Database • IBM DB2 Universal Database™ • Microsoft SQL Server • Oracle	**Terminal Servers** • Citrix (64-bit versions) • Microsoft Terminal Server (64-bit versions)	
ERP / CRM • PeopleSoft • SAP • Siebel		
Java Servers • BEA WebLogic • IBM WebSphere		
Virtual Machine Servers • Microsoft Virtual Server • VMware ESX Server		

Table 1. *Dual-core benefits for servers by application category*

System Overview

The IBM systems used in this paper to compare Intel's single-core processor to dual-core processor are the x346 and the new x3650 (*Figure 1* shows a side-by-side diagram of the system architectures of the x346 and x3650). In addition to the transition from single-core to dual-core, the Xeon Processor 51xx series introduces Intel's new Core™ Microarchitecture[3]. Because Woodcrest is a new dual-core processor, the cache structure differs from the earlier single-core Xeon processors (codenamed Irwindale). Woodcrest has two L1 caches (one dedicated per core) and shares the L2 cache across both cores.

Apart from the different processor cores, there are other notable changes that greatly enhance the benefits of the x3650 over the x346. By effectively doubling the number of processors from two to four, the other system components, such as FSB bandwidth, memory bandwidth, and I/O bandwidth, had to be scaled to handle the increase in processing power. To process the increase in traffic without increased congestion, the x3650 more than doubles the FSB bandwidth versus the x346—by a factor of 2.6X. In addition, memory bandwidth is scaled by 3.3X to support a greater number of memory DIMMs, and the x3650 also has a 2.5X increase in its I/O slot bandwidth. All these changes were designed with the capacity to support future quad-core systems.

---

[3] For more details about Core Microarchitecture: http://www.intel.com/technology/architecture/coremicro/index.htm

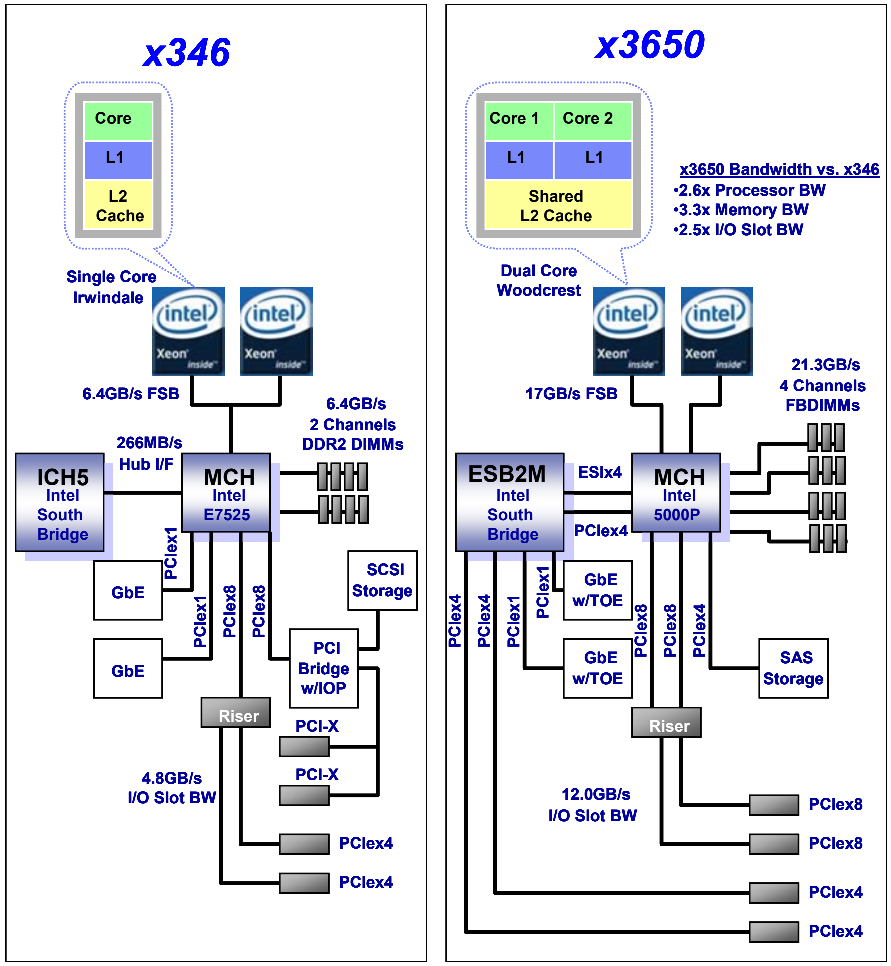

Figure 1. *x346 vs. x3650*

The experiments and results discussed throughout this paper are based on the systems above. Although the clock frequencies in some of the experiments may differ, the underlying schematics are similar.

Performance Comparisons

To help understand the application performance and performance-per-watt differences between single-core (SC) and dual-core (DC) processors, IBM performed a number of industry benchmarks to compare the x346 using single-core (SC) Xeon processors and the x3650 using dual-core (DC) Xeon processors. These benchmarks can be useful for comparing processors under varying workloads.

Overall Performance Improvements

To compare the x346 and x3650 in overall performance, two benchmarks are used: Microsoft Exchange Server MMB3 and SPECweb2005. Exchange Server MMB3 is an industry-standard benchmark used to measure a server's capability to run e-mail-based operations. SPECweb2005 simulates sending requests to a web server through a broadband internet connection.

While dual-core processors attract much attention as the source of performance improvement, *all* changes in the configurations of the servers can have an impact on the performance results. The main differences between the two servers used in this experiment are listed in *Table 2*[4]. To allow for an increase in performance, the x3650 significantly increased the total memory and disk drive capacities. This helps accommodate the increase in workload ability that comes along with the dual-core processor. Technology has also driven a change to better performing components. The switch to a new generation of memory technology and the SAS drives are each key contributing factors that allow the IBM dual-core systems to outperform previous single-core systems.

	x346	x3650
Processors / Cores	2 x **3.8GHz SC Xeon** (Irwindale) / **2** cores	2 x **3.0GHz DC Xeon** (Woodcrest) / **4** cores
Memory	**4GB PC2-3200** DDR2 **Registered DIMMs** (Exchange); **16GB PC2-3200** DDR2 **Registered DIMMs** (SPECweb2005)	**8GB PC2-5300** DDR2 **Fully Buffered DIMMs** (Exchange); **24GB PC2-5300** DDR2 **Fully Buffered DIMMs** (SPECweb2005)
Hard Disk Drives	**73GB** 15K RPM **Ultra320 SCSI**, **9,180GB** 15K RPM Fibre Channel (Exchange); **3,296GB** 15K RPM **Ultra320 SCSI** (SPECweb2005)	**876GB** 15K RPM **SAS**, **12,348GB** 15K RPM Fibre Channel (Exchange); **146GB** 15K RPM SAS, **1,260GB** 15K RPM Fibre Channel (SPECweb2005)
Web Server Software (SPECweb2005 only)	Zeus Web Server v4.2r4 (x86-64)	Accoria Rock Web Server v1.3.3 (x86-64)

Table 2. *System configuration for Exchange and SPECweb2005 benchmarks*

In order to show the benefits of the x3650, *Figure 2* has been normalized to the x346 scores. (In other words, all results are expressed as a percentage of the x346 scores.) Exchange shows a **34%** increase in performance for the x3650 over the x346, while SPECweb2005 exhibits a **111%** increase. These dramatic increases in performance are directly related to the component improvements as described in the previous section: increasing the processor, memory, and I/O bandwidths, as well as other differences between the x346 and x3650.

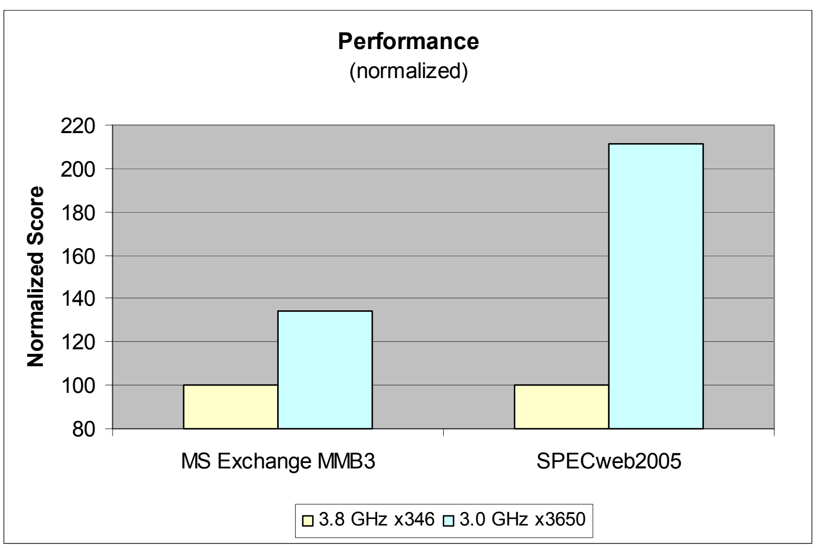

Figure 2. *Performance comparison of x346 and x3650*

[4] For full disclosures, go to http://www.microsoft.com/exchange/evaluation/performance/default.mspx and http://www.spec.org/web2005/results/

Final performance note: To see the full benefits of dual-core processors, especially with 64-bit software, it is necessary to provide ample memory for the software. If previously you supplied 2GB of memory per single-core processor, you should consider doubling those numbers when using dual-core processors. (In-house application-specific testing is advisable in any case to determine actual needs.)

Performance Per Watt

In the past, as workloads increased, infrastructure responded to meet those needs. As infrastructure grew, so did power consumption. With DC servers, this is not the case. From the previous discussion, it is clear that DC servers can reduce the size of your overall infrastructure purely in terms of performance. In this section, you will see that power consumption of dual-core servers is reduced, out-performing single-core servers.

The server configurations used to run the benchmarks for the performance-per-watt measurements are listed in *Table 3*. Some important differences to note are the move to a new generation of memory technology and SAS disk drives. Both of these technologies offer increased performance.

	x346	**x3650**
Processors / Cores	2 x **3.6GHz SC** Xeon (Irwindale) / **2** cores	2 x **2.33GHz DC** Xeon (Woodcrest) / **4** cores
Memory	4 x 1GB (**PC2-3200** DDR2 **Registered DIMMs**)	4 x 1GB (**PC2-5300** DDR2 **Fully Buffered DIMMs**)
Hard Disk Drives	**6** x 73GB **15K** RPM **3.5**-inch **SCSI**	**8** x 73GB **10K** RPM **2.5**-inch **SAS**
Disk Configuration	**6**-Drive RAID-5	**8**-Drive RAID-5
NICs	2 x 1Gb (planar)	2 x 1Gb (planar)
RAID Controller	ServeRAID-**7k** (SCSI)	ServeRAID-**8k** (SAS)

Table 3. *System configuration for file server and web server benchmark*

The performance-per-watt metric measures a server's effectiveness in maximizing performance for every watt of power consumed. To arrive at a score for this metric, the performance score is divided by the peak power measured during the maximum throughput of the server. This number then represents the performance-per-watt measurement and the results are normalized as before.

Figure 3 shows the performance-per-watt results of a file server benchmark and a web server benchmark run on the x346 and x3650. The file server benchmark measures how well a file server handles I/O requests, while the web server benchmark measures how well a server processes Java web requests. The file server benchmark shows a **90%** increase for the x3650 server over the x346 and the web server benchmark a **175%** increase for the x3650 server. These are dramatic increases in performance-per-watt that can not be attributed to one specific difference between the x346 and x3650. Especially for the file server benchmark, the new SAS disk drives play a major role in the results, due to the heavy I/O traffic simulated.

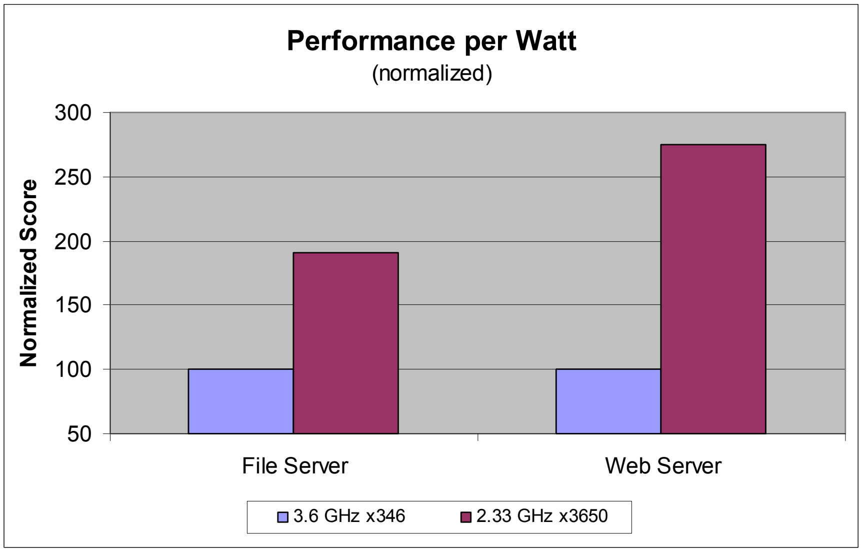

Figure 3. *Performance-per-watt comparison of IBM servers using SC and DC Intel processors*

Conclusion

In yesterday's marketplace, the demand for higher performance was met by decreasing transistor size and increasing clock frequency. Today, technology has moved in the direction of increasing performance by adding processor cores. The x3650 shows performance increases over the x346 of up to **34%** on the Microsoft Exchange benchmark and **111%** on the SPECweb2005 benchmark.

Even with the performance boosts, the real benefits of dual-core systems are in power usage. There the x3650 shows performance-per-watt increases over the x346 of up to **90%** on the file serving benchmark and **175%** on the web serving benchmark. *Figures 4* and *5* summarize the potential of an x3650 infrastructure in a scenario where the existing infrastructure uses **21 x346** servers. By comparison, the same amount of work could be done with only **11 x3650** servers, using nearly *two-thirds less* power.

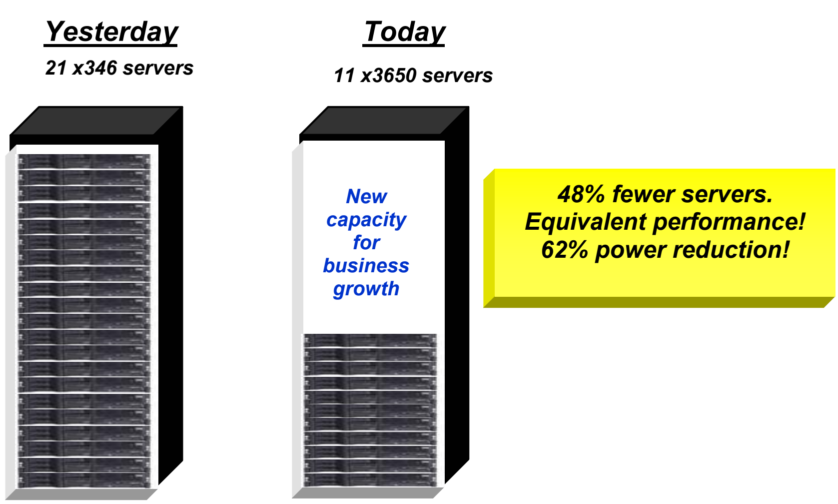

Figure 4. *IBM x346 vs. IBM x3650: fewer servers, less power, equal performance*

Alternatively, using the *same* number of x3650 servers, nearly *twice* the work can be performed, while still reducing the power draw by *more than a quarter*.

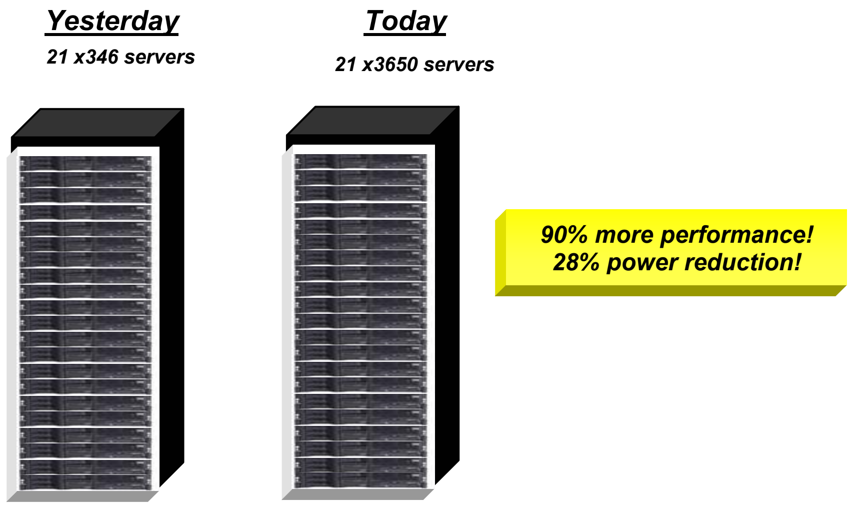

Figure 2. x346 vs. x3650: more performance, less power

The age of dual-core processing has arrived and IBM is ready to deliver with a complete line of server solutions (3U x3950 and x3850, 2U x3650, and 1U x3550, as well as x3800 and x3500 high-end towers, and x3400 entry tower) to meet the needs to today's marketplace. These IBM servers offer more performance and better power management than ever before. If you have been searching for more capacity for new business growth, more performance per watt of power, and the ability to reduce overall infrastructure size, now is the time and these are the systems.

Acknowledgments

Special thanks to the IBM System x Performance and Analysis group, especially Matthew Eckl and Sylvester Cash, and all who contributed during the reviewing process.

For More Information

IBM System x and xSeries Servers	**ibm.com**/systems/x
IBM Rack Configurator	**ibm.com**/pc/us/eserver/xseries/library/configtools.html
IBM Configuration and Options Guide	**ibm.com**/servers/eserver/xseries/cog

Legal Information

IBM Systems and Technology Group
Dept. U2SA
3039 Cornwallis Road
Research Triangle Park, NC 27709

Produced in the USA
June 2006

For a copy of applicable product warranties, write to: Warranty Information, P.O. Box 12195, RTP, NC 27709, Attn: Dept. JDJA/B203. IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven or ClusterProven. Telephone support may be subject to additional charges. For onsite labor, IBM will attempt to diagnose and resolve the problem remotely before sending a technician.

IBM, the IBM logo, the e-business logo, PowerExecutive, System x, xSeries and Xtended Design Architecture are trademarks of IBM Corporation in the United States and/or other countries. For a list of additional IBM trademarks, please see http://**ibm.com**/legal/copytrade.shtml.

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

Linux is a registered trademark of Linus Torvalds.

Microsoft, Windows and the Windows logo are trademarks or registered trademarks of Microsoft Corporation.

Other company, product and service names may be trademarks or service marks of others.

IBM reserves the right to change specifications or other product information without notice. References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. IBM PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This publication may contain links to third party sites that are not under the control of or maintained by IBM. Access to any such third party site is at the user's own risk and IBM is not responsible for the accuracy or reliability of any information, data, opinions, advice or statements made on these sites. IBM provides these links merely as a convenience and the inclusion of such links does not imply an endorsement.

Information in this presentation concerning non-IBM products was obtained from the suppliers of these products, published announcement material or other publicly available sources. IBM has not tested these products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

MB, GB and TB = 1,000,000, 1,000,000,000 and 1,000,000,000,000 bytes, respectively, when referring to storage capacity. Accessible capacity is less; up to 3GB is used in service partition. Actual storage capacity will vary based upon many factors and may be less than stated.

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will depend on considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

Maximum internal hard disk and memory capacities may require the replacement of any standard hard drives and/or memory and the population of all hard disk bays and memory slots with the largest currently supported drives available. When referring to variable speed CD-ROMs, CD-Rs, CD-RWs and DVDs, actual playback speed will vary and is often less than the maximum possible.

XSW02525-USEN-00