Supercharging your high-performance computing
June 2005

The Benefits of Dual-Core Processors in High-Performance Computing

By Mark T. Chapman
IBM Systems and Technology Group

Executive Overview

This paper examines the real-world advantages of dual-core processors; specifically how the new dual-core AMD™ Opteron™ processors benefit users in a high-performance computing (HPC) environment.

The next step in advanced processor technology, following the ongoing transition to 64-bit computing, is the introduction of dual-core processors. Typical processors chips contain one processor core (the “brains” of the processor) surrounded by supporting circuitry, such as on-chip L2 cache, a memory controller (in some designs), and so on. 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 dual L2 caches and the other supporting circuitry. This offers the potential for greater performance and reduced latency, with lower power draw and heat output, than two physical processors would provide.

Contents

Dual-Core Benefits

There are two major advantages gained by switching from single-core processors to dual-core processors: greater performance and lower electrical costs.

Anyone who follows the IT industry knows that it is increasingly difficult to scale up the clock rate of processors—they become more difficult and more expensive to design and build. The higher the clock frequency, the more waste heat is generated. The hotter the chip, the hotter the system runs and the more challenging it is to keep everything cool. Raising clock frequency as a means of increasing system performance has reached a point of rapidly diminishing returns. Other, more effective, approaches are needed.

To date, the simplest alternative has been to increase the size of the internal L2 cache and perhaps add an external L3 cache. This allows 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. Still more innovation is required.

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” your 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 (virtual) processors. Theoretically, it can process two threads indefinitely. In reality, this rarely happens because those virtual 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 resource. In addition, if the software stack is largely single-threaded, the second virtual processor may be idle much of the time, rendering its value moot.

By contrast, with a dual-core dual-threaded processor, you actually have two physical processors on 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 multithreaded 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 2.0GHz *dual*-core Opteron processor offers significantly more total throughput than a 2.6GHz *single*-core processor. However; a uniprocessor server containing *one* 2.0GHz *dual*-core chip will *not* be as fast as a two-socket server utilizing *two* 2.6GHz *single*-core chips. Each of the single cores in the two-socket server runs faster than the individual cores in the dual-core uniprocessor server. In performance terms, it makes more sense to think of a server containing *one dual*-core processor as a performance upgrade to a server with *one single*-core processor, not as an alternative to a server with *two single*-core processors.

The same is true when comparing a server with two dual-core processors with a conventional 4-way single-core server. Due to the significantly lower cost of a dual-socket/dual-core server when compared with a 4-way server, it becomes tempting to switch, especially for low-end 4-way uses. Before switching, however, consider the other distinctions between a 4-way server and a 2-way server: 4-way servers typically are physically larger than 2-way systems, which allow them to hold more memory, more drives and more I/O slots. Consider whether the two-socket/dual-core server offers enough expandability for your "4-way" needs. While the processor is a *sine qua non*, it is only a *piece* of the total server puzzle.

Another advantage of dual-core processors is the power and thermal reduction. One *dual*-core chip offers nearly the same performance as two *single*-core chips of the same clock rate, while using half the power and producing half the waste 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 processors*, rather than the number of cores; therefore the transition from single-core to dual-core processors should have little or no effect on software licensing fees for most customers.

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 I/O-intensive applications.

Will You Benefit?

There are a number of categories of applications and middleware that can benefit from dual-core processors. To help you determine whether these new processors will help you, the following tables summarize the categories and representative applications that should see significant benefit, moderate benefit, or little benefit from a switch from single-core to dual-core processors. *Table 1* shows the benefits for dual-core servers, while *Table 2* shows the benefits for workstations. (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 simply listed in alphabetical order.)

Significant	Moderate	Little
HPC Servers • Automotive • Aeronautical • EDA • Geophysical • Life sciences	**Java Servers** • BEA WebLogic • IBM WebSphere	**Terminal Servers** • Citrix • Microsoft Terminal Server
DCC Servers • Pixar RenderMan	**Web Servers** • Apache • Microsoft Internet Information	**E-Mail Servers** • Microsoft Exchange • Sendmail

Significant	Moderate	Little
	Server	
Data Mining • MicroStrategy • SAS	**Virtual Machine Servers** • Microsoft Virtual Server • VMware ESX Server	
Database • IBM DB2 Universal Database™ • Microsoft SQL Server • Oracle	**File Servers** • CFIS • NFS • Samba	
ERP / CRM • PeopleSoft • SAP • Siebel		

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

Significant	Moderate	Little
DCC Servers • 3D Studio Max • Alias Maya • SoftImage XSI	**Video Editing** • Avid • Pinnacle	**CAD** • Autodesk AutoCAD • EDS Unigraphics • PTC Pro/ENGINEER • SolidWorks
CAE Servers • Algor Nastran • ANSYS • Fluent	**Software Development** • Compilers • Java engines	
EDA • Cadence • Mentor Graphics • Synopsis		
Scientific • Life sciences • Oil and gas		

Table 2. *Dual-core benefits for **workstations**, by application category*

Opteron, an HPC Powerhouse

In an AMD Opteron™ processor-based server, memory has its own high-speed path to the processor, rather than having to share a front-side bus with other devices. This gives Opteron a high-speed link to memory—memory is accessed at the full clock rate. Moreover, because each Opteron contains its own integrated memory controller, installing a second Opteron means having *two* high-speed data paths to memory, effectively providing up to double the throughput. (Having an integrated memory controller also reduces memory latency, for faster memory access.) Single-core Opteron processors are currently offered in 2.2, 2.4 and 2.6GHz clock speeds (Models 248, 250 and 252 respectively).

Although the dual-core Opteron processors (Models 265, 270 and 275) run at “only” 1.8, 2.0 and 2.2GHz, respectively, the second core more than makes up the difference in clock rate. One 2.0GHz (270) dual-core chip is considerably faster on typical HPC benchmarks than one 2.6GHz (252) single-core chip. While one *dual*-core Opteron may not be as fast as two

single-core Opteron processors running at a higher clock rate, two dual-core processors *will* be considerably faster than two single-core processors[1]. With the AMD architecture and dual-core computing, clock rate becomes less meaningful for judging performance.

The dual-core design offers specific benefits to software. Single-core Opteron processors were already outstanding when it came to memory bandwidth-constrained applications, offering massive memory bandwidth. Doubling the number of cores means that the amount of on-chip cache and the number of processor registers doubles as well. This helps considerably in core-constrained applications. (I/O intensive applications are largely unaffected by the presence of dual cores.) In addition, AMD provides power management features (called *PowerNow!™ technology with Optimized Power Management*) in its Opteron processors that allow the processor to automatically step down power consumption when not fully utilized. PowerNow! technology can reduce CPU power at idle by up to 75%. IBM plans to support this technology once operating system support is available.

The single- and dual-core "flavors" of Opteron processors are compatible with 32-bit and 64-bit versions of both Microsoft® Windows® 2003 and Linux®, offering great flexibility in terms of OS support. (Look for operating systems that specify AMD Opteron, AMD64 or x86-64 support to take full advantage of the processor's features.) Opteron extends the original Intel x86 core used in 32-bit Intel processors. Opteron can natively run both 32-bit and 64-bit x86 software, concurrently, at extremely high levels of performance.

System Comparisons

Because one size *doesn't* fit all, IBM offers a selection of systems designed around Opteron processors. IBM's portfolio of single- and dual-core enabled servers and workstations currently consists of the *IBM @server 326* (e326), the *AMD Opteron LS20 for the IBM @server BladeCenter™* (LS20) and the *IBM IntelliStation® A Pro*. *Table 1* compares the major features of each single-core (SC) or dual-core (DC) Opteron processor-based IBM system, to give you a better idea of which best fits your specific needs. (Significant distinctions are highlighted in bold.)

Features	e326	LS20	A Pro
Form factor	**1U rack-optimized server**	**30mm blade server**	**Minitower (4U) workstation**
Opteron processor models supported	**242, 244,** 246, **248,** 250, 252 (SC); **265,** 270, 275 (DC)	246, 250, 252 (SC); 270 (DC)	250, 252 (SC); 275 (DC)
Number of processor sockets	2	2	2
Number of processor cores per socket	1 or 2	1 or 2	1 or 2
Maximum processor **sockets** per 42U rack	84	**168**	20
Maximum processor **cores** per 42U rack	168 (dual-core); 84 (single-core)	**336** (dual-core), 168 (single-core)	40 (dual-core); 20 (single-core)
Maximum clock rate[2]	2.6GHz (Model 252, single-core); 2.2GHz (Model 275, dual-core)	2.6GHz (Model 252,single-core); 2.0GHz (Model 270, dual-core)	2.6GHz (Model 252, single-core); 2.2GHz (Model 275, dual-core)

[1] Visit www.amd.com/opteronperformance to learn more about performance benchmarks for AMD processors.

[2] Maximum clock rate of processors offered for these models as of March 31, 2005. Some submodels may use processors running at a lower clock rate. Future processors may operate at faster frequencies.

Features	e326	LS20	A Pro
Thermal design power (maximum) per proc.	95W (single-core/dual-core)	**68W** (single-core/dual-core)	95W (single-core/dual-core)
# of Memory sockets	8	**4**	8
Type of memory supported	400MHz DDR (PC3200)	400MHz DDR (PC3200)	400MHz DDR (PC3200)
Uni / 2-socket memory bus bandwidth	6.4GB/sec / 12.8GB/sec	6.4GB/sec / 12.8GB/sec	6.4GB/sec / 12.8GB/sec
System memory capacity[3]	16GB	**8GB** / 16GB[4]	16GB
HDD form factor	3.5"	**2.5"**	3.5"
HDD technologies supported	Hot-swap Ultra320 SCSI; Fixed SATA	Fixed Ultra320 SCSI	Hot-swap Ultra320 SCSI; Fixed SATA
Maximum HDD capacities supported	**300GB** SCSI; 160GB SATA	73.4GB SCSI	146.8GB SCSI; **250GB** SATA
# of HDD bays	2	2	**4**
Maximum internal HDD capacity	**600GB** SCSI; 320GB SATA	146.8GB SCSI	587.2GB SCSI; **500GB** SATA
HDD RPMs	10K (**300GB**) / 15K (**146.8GB**) — SCSI; 7,200 — SATA	10K — SCSI	10K (**146.8GB**) / 15K (**73.4GB**) — SCSI; 7,200 — SATA
I/O slots available	(2) PCI-X (one adapter at 133MHz or two at 100MHz)	(1) legacy PCI-X or (2) SFF[5] PCI-X 133MHz daughtercards	(1) 133MHz PCI-X; (**4**) 100MHz PCI-X
Gigabit Ethernet ports	2	2 (via BladeCenter chassis)	**1**
Internal/external switches	External	**Internal** to BladeCenter chassis	External
USB ports	(2) USB 1.1	(2) USB 1.1	**(5)** USB **2.0**
IEEE 1394 (FireWire) ports	0	0	**2**
32-bit OS support	Windows/Linux/ **NetWare**	Windows/Linux/ **Netware**	Windows
64-bit OS support	Windows/Linux	Windows/Linux	Linux
64-bit OS availability	• MS Windows Server 2003 x86-64 • Novell SLES 8 AMD-64 (post-GA) • Novell SLES 9 x86-64 • RHEL 3 x86-64 • RHEL 4 x86-64 • VMware ESX Server 3	• MS Windows Server 2003 x86-64 • MS Windows XP Professional (post-GA) • Novell SLES 9 x86-64 • Novell Open Enterprise Server (NetWare) 6.5 (post-GA) • RHEL 3 x86-64	• RHEL 3 Workstation

[3] Although not a function of the processor architecture, maximum system memory capacity can affect your decision of which server is best suited to your needs.

[4] When 4GB DDR2 DIMMs are available.

[5] Small Form Factor.

Features	e326	LS20	A Pro
		• RHEL 4 x86-64 • VMware ESX Server 3	

Table 3. *64-bit processor comparison for 2-socket servers*

Performance Comparisons

In order to help you decide between single-core (SC) and dual-core (DC) processors, and between the e326 and the LS20, IBM performed a number of industry benchmarks to compare the fastest single-and dual-core Opteron processors available for the e326 and the LS20[6]. These benchmarks can be useful for comparing processor performance under varying workloads.

Linpack

Graph 1 shows the results of the Linpack benchmark, which is used to solve a dense system of linear equations. Linpack typically benefits from faster processor clock rates.

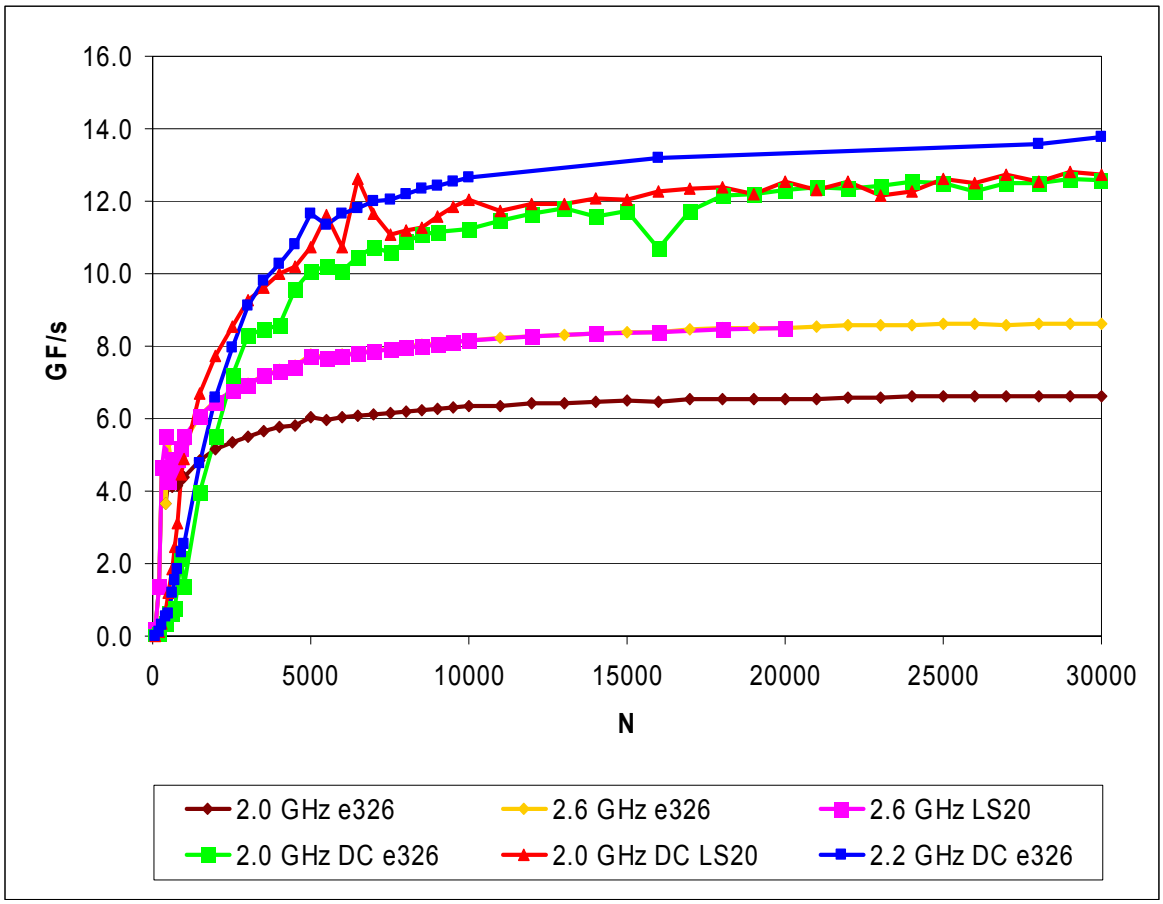

Graph 1. *Comparison of IBM servers using single-core and dual-core Opteron processors*

As you can see from the graph, the dual-core processor-equipped servers performed significantly better on the Linpack benchmark than the single-core processor models. Even though the clock rates of the individual cores in the DC processors were slower than their single-core counterparts, there were twice as many cores doing the work. Comparing an e326 using the fastest available (2.2GHz) *DC* Opteron processor (Model 275) to an e326 using the fastest (2.6GHz) *SC* Opteron

[6] The 2.6GHz single-core LS20 server included 4GB of RAM. All other configurations included 8GB of memory.

processors (252) shows a **50%-60%** improvement at most benchmark levels. Comparing like-versus-like, an e326 using two 2.0GHz *DC* Opteron processors (265) was more than **90%** faster at most levels than an e326 with two 2.0GHz *SC* Opteron processors (246), and more than *twice as fast* at some levels.

SPEC CPU2000

The SPEC CPU2000 benchmarks measure the performance of a system's processors and memory. *Table 4* compares published[7] SPECint_rate2000 and SPECfp_rate2000 scores for single-core LS20 servers versus dual-core LS20 and e326 servers.

Server/Processor	SPECint_rate2000	SPECfp_rate2000
LS20 with two SC Opteron 246 (**2.0GHz**, **1MB** L2 cache)	30.7	34.8
LS20 with two SC Opteron 252 (**2.6GHz**, **1MB** L2 cache)	40.7	46.5
LS20 with two *DC* Opteron 270 (**2.0GHz**, **2 x 1MB** L2 cache)	60.1	61.9
e326 with two *DC* Opteron 275 (**2.2GHz**, **2 x 1MB** L2 cache)	64.0	65.9

Table 4. *CPU2000 benchmark results for LS20 and e326 servers*

Note that in the SPECint_rate2000 benchmark (which measures **integer** performance) the e326 equipped with two 2.2GHz Opteron 275 processors is more than **57%** faster than the LS20 with two 2.6GHz Opteron 252 processors. In the SPECfp_rate2000 benchmark, which measures **floating-point** performance, the e326 with *DC* 2.2GHz Opteron 275 processors is nearly **42%** faster than the LS20 with two *SC* 2.6GHz Opteron 252 processors. This clearly shows that on suitable workloads the fastest dual-core system is much faster than the fastest single-core system, even though the single-core system has a higher clock frequency.

Comparing the LS20 equipped with two 2.0GHz *DC* Opteron 270 processors versus an LS20 equipped with two 2.0GHz *SC* Opteron 246 processors, the *DC* system is nearly **96%** faster in integer performance, and **78%** faster in floating-point performance at the same clock rate.

It is worth noting that for SPECint_rate2000 the 2.0GHz DC system is almost exactly twice as fast as the 2.0GHz SC system. This is because this benchmark reflects a workload focused almost entirely on the processor core. The SPECfp_rate2000 results also show excellent scaling, though not quite as impressive as the SPECint_rate2000 results. The difference occurs because the SPEC CFP2000 Rates benchmark (which uses the SPECfp_rate2000 metric) places much higher demands on memory bandwidth than does the SPEC CINT2000 Rates benchmark (which uses the SPECint_rate2000 metric). Note that the *memory bandwidth* is proportional to the number of processor *sockets* in use, whereas the *processing power* is proportional to the number of processor *cores* in use. By using dual-core processors, we increase the amount of processing power, but not the bandwidth to memory. Thus, SPECfp_rate2000 does not scale quite as well.

STREAM

Another industry benchmark important to those involved in high-performance computing is STREAM. The STREAM benchmark is a simple synthetic benchmark program that measures sustainable memory bandwidth (in MB/s) and the corresponding computation rate for simple vector kernels. The Opteron platform supports high bandwidth-to-memory, allowing it to excel at the type of processing that emphasizes memory bandwidth rather than cache size or clock frequency. In a dual-core system, the number of processor cores increase but not the overall

[7] Results available at http://www.spec.org.

memory bandwidth. Nevertheless, through increased efficiency in the dual-core processors and advances in IBM server technology, memory *throughput* does increase somewhat. The most dramatic results can be seen in the LS20, where enhancements in the LS20 design allow the **2.0GHz** dual-core model to achieve higher scores than either the **2.0GHz** dual-core e326 or the **2.2GHz** DC e326 model. *Table 5* compares results for single-core e326 and LS20 servers versus dual-core servers.

Server/Processor	STREAM Copy/Scale	STREAM Add/Triad
LS20 with two SC Opteron 246 (**2.0GHz**, **1MB** L2 cache)	8,535	8,340
e326 with two SC Opteron 252 (**2.6GHz**, **1MB** L2 cache)	8,273	8,414
LS20 with two SC Opteron 252 (**2.6GHz**, **1MB** L2 cache)	8,760	8,550
e326 with two *DC* Opteron 270 (**2.0GHz**, **2 x 1MB** L2 cache)	9,093	9,144
LS20 with two *DC* Opteron 270 (**2.0GHz**, **2 x 1MB** L2 cache)	**10,670**	**9,601**
e326 with two *DC* Opteron 275 (**2.2GHz**, **2 x 1MB** L2 cache)	9,181	9,284

Table 5. *STREAM benchmark results for LS20 and e326 servers*

Using the **Copy** and **Scale** tests, the e326 equipped with two 2.2GHz *DC* processors is approximately **5%** faster than the e326 with two 2.6GHz *SC* processors. Similarly, the e326 with 2.0GHz *DC* processors is more than **6%** faster than the 2.0GHz *SC* e326 system. Comparing the LS20 equipped with *DC* 2.0GHz processors to the LS20 with 2.6GHz *SC* processors shows a nearly **22%** improvement. Compared to the 2.0GHz *SC* LS20 system, the 2.0GHz *dual*-core system is **25%** faster.

Using the **Add** and **Triad** tests**,** the e326 with *DC* 2.2GHz processors is more than **8%** faster than the e326 with two *SC* 2.6GHz processors. Similarly, the e326 with 2.0GHz *DC* processors is nearly **10%** faster than the 2.0GHz *single*-core e326 system. The LS20 with *DC* 2.0GHz processors is more than **12%** faster than the LS20 with two *SC* 2.6GHz processors. Comparing the LS20 equipped with two 2.0GHz *DC* processors to an LS20 equipped with two 2.0GHz S*C* processors, the dual-core system is **15%** faster at the same clock rate.

Final performance note: A server with one DC processor is much like a server with two SC processors in one respect. In order 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 for a single processor and 4GB in a 2-way single-core server, you should consider doubling those numbers when using dual-core processors. (In-house application-specific testing is advisable in any case to determine your actual needs.)

Given the performance advantages of dual-core processors, does this mean the end of single-core processors? Eventually, yes; however for now, there are still advantages to using single-core processors:

1) If you are using a single-threaded application that would benefit from the fastest possible clock rate, a server with a 2.4 or 2.6GHz *single*-core Opteron processor would probably run the application faster than one with a 2.2GHz *dual*-core Opteron.
2) Dual-core-enabled servers can offer significant advantages in price/performance over single-core systems. However, if your budget is tight, a server containing a 2.0GHz *single*-core Opteron, for example, may cost significantly less than one using a 1.8GHz *dual*-core Opteron in the short term.

LS20 or e326 Server?

So which server is the better fit for your HPC requirements? It all depends on your needs.

Price/Performance

When deployed in large numbers, LS20 blades combined with a BladeCenter chassis (*Figure 1*) can offer tremendous cost savings in the following areas:

- Simplified centralized systems management
- Integration and consolidation of switches, fans, power supplies, optical drives and other components (which reduces the number of internal components to service)
- Lower ongoing costs for data center floor space, power and cooling
- Fewer cables to purchase and manage

In addition, the LS20 uses special low-voltage (68 watt) versions of the Opteron processors (both single-core and dual-core). They provide the same performance as the typical higher-voltage processors, but use about **28%** less power. This can save you significant money over time on power and cooling costs.

Conversely, the 1U e326 can be the more cost-effective choice when deployed in small numbers or when systems management and data center square footage are not a top priority.

Figure 1. 7U *BladeCenter chassis with 14 LS20 blades*

Despite using many of the same processors, there are performance differences between the two servers. The e326 design provides a faster communications interconnect than the LS20 and twice the memory capacity (at least until 4GB DIMMs are available for the LS20), as well as larger, faster internal hard disk drives. Go with LS20 when memory throughput is of greater importance.

Flexibility / Integration

The e326 (*Figure 2*) offers more flexibility in internal configuration than the LS20, by supporting a choice of SATA or hot-swap SCSI models, compared with only fixed SCSI in the LS20. Of course, if you prefer to use external SAN storage, the speed and capacity of internal drives would be irrelevant. The e326 also supports two conventional PCI-X adapters, vs. one or two PCI-X daughtercards in the LS20.

Due to its integrated chassis design, BladeCenter offers LS20 users the flexibility of using a number of choices of internal communications switches interchangeably and without the bother of major recabling whenever switches are added or removed. The integration also provides higher server density than the e326. LS20 also offers maximum availability due to the extensive use of hot-swap and redundant components in BladeCenter.

Figure 2. *e326 1U server*

In general, choose e326 for low-cost, high-performance cluster nodes and front-end applications that do not require high availability. Choose LS20 when your data center is space-constrained or you require extensive infrastructure integration.

IntelliStation A Pro?

For single use, or small clusters, the standalone IntelliStation A Pro (*Figure 3*) requires none of the rack and data center infrastructure of the servers. If your needs are for large internal storage capacities or many adapters, the IntelliStation A Pro offers considerable expandability, with six HDD bays and five available PCI-X adapter slots. For high-performance graphics needs, the A Pro offers a variety of high-powered adapters, ranging from 2D up through extreme 3D. Choose IntelliStation A Pro when a stand-alone system or leading-edge graphics is your top priority.

Figure 3. *IntelliStation A Pro*

Conclusion

The new Opteron processors combine the advantages not only of 64-bit processing but also dual-core processing. Moreover, the dual-core Opteron processors are completely compatible with their single-core brethren. As you have seen, dual-core processors provide significant performance advantages over single-core processors for high-performance computing tasks. Heavily multithreaded applications will potentially see major benefits from dual-core processors by having up to eight threads running at once in a 2-socket server (two threads per core). Even single-threaded applications will gain by having the functional equivalent of four processors available in a "2-way" server. However, to see the maximum performance benefit, it is advantageous to increase the amount memory available to the applications, just as if you were doubling the number of physical processors in the server.

The age of dual-core Intel-compatible computing has arrived. To turbocharge your HPC operations, include individual IntelliStation A Pro workstations and e326 and LS20 servers featuring dual-core Opteron processors—or run groups of single- and dual-core servers in an IBM @server Cluster 1350 configuration.

Appendix A: System Summaries

The following is additional background information about the e326, the LS20, the IntelliStation A Pro and the IBM @server Cluster 1350.

IBM @server 326

Advantages:

- *High memory density*
- *Low deployment cost*

For those users seeking a **1U rack-optimized** server, IBM offers a choice of single-core and dual-core Opteron processors in the 2-socket **e326** server. The e326 is optimized for HPC use, offering the features necessary in high-performance computing without burdening the server with features seldom used in that environment.

Beyond the components typically found in a 1U server, the e326 takes advantage of IBM's Calibrated Vectored Cooling™ design methodology (which provides extremely efficient server cooling) to offer features *not* generally found in this class of server, such as *eight* DIMM slots (providing up to 16GB of high-speed memory using inexpensive 2GB DIMMs)—instead of the usual six.

AMD Opteron LS20 for the IBM @server BladeCenter

Advantages:

- *Integrated infrastructure*
- *High processor density*
- *Lowest long-term operational costs*

IBM's newest blade server, the LS20, is similar in most respects to its sibling two-socket blade servers: the Intel Xeon® processor-based HS20 and the IBM PowerPC® 970FX processor-based JS20. The LS20 uses the same BladeCenter chassis as the other blade servers (and can be mixed-and-matched with them), as well as most of the same options.

As with IBM's other blade servers, the LS20 leverages the tremendous integration and consolidation of the BladeCenter chassis. By moving the power supplies, fans, CD-ROM and floppy drives, and the systems management controllers from the individual servers to the chassis, IBM was able to achieve twice the server rack density of traditional 1U servers. Further *rack* density was achieved by moving the Ethernet, Fibre Channel, Myrinet and other switches from the rack into the BladeCenter chassis. One 7U chassis holds up to 14 servers plus *four* switch modules. Connecting the servers to the chassis using a midplane allowed IBM to eliminate all of the Ethernet-to-switch cables, for example.

Instead of hundreds of power, communication, systems management and KVM cables per rack, a BladeCenter solution reduces the clutter to a manageable minimum. To illustrate, here is a sample configuration, comparing the rack and cabling requirements for **84** IBM blade servers (in **six** BladeCenter chassis), versus **84** typical 1U servers. Both examples include full redundancy.

Feature	BladeCenter	1U Server
Number of servers *per 42U rack*	84	42
Number of racks required for 84 servers	1	2
Number of power cords	24	168
Number of systems management cables	12	168
Number of Ethernet cables[8] from servers to switches	0[9]	168
Number of Ethernet cables exiting the switches	Up to 24	Variable[10]
Number of KVM cables	6	168
Total number of cables in rack	Up to **66**	**672+**

Table 6. *Cable comparison between BladeCenter and 1U servers*

With the same number of 1U servers as blade servers, you would need *twice* the data center square footage, incur *twice* the rack hardware cost and require at least *ten times* as many cables. The cost savings in reduced administrative time can be considerable as well. Moreover, due to the widespread use of hot-swap and redundant components, the entire BladeCenter chassis can be configured for *no single point of failure*. With up to four switches, four power supplies, four blowers, two management controllers and a redundant midplane, BladeCenter offers redundancy for everything, with provisions for automatic failover. For even higher redundancy, you can also cluster multiple blades.

[8] Using two Gigabit Ethernet controllers per server.

[9] All Ethernet connections between blade servers and switches are routed through the BladeCenter chassis midplane. No cables are required.

[10] Depending on configuration.

A standard 42U rack can hold up to 84 two-processor LS20 blade servers (or any combination of 84 LS20, HS20 and JS20 blades). Using single-core processors, this means a rack can support up to **168** high-performance processors and **24** communications switches. Using dual-core processors, the number of processor cores per rack doubles to **336**, providing an incredible amount of processing power in a single rack.

IBM @server Cluster 1350

Advantages:

- *Turnkey clustering solution*
- *Clusters optimized for HPC*

Scientists, engineers and researchers have long chosen clustered servers as the preferred solution for large, complex, computationally intensive problems. Clusters have also proven to be a cost-effective method for managing mixed HPC workloads. It is common today for corporations, universities and government labs to utilize server clusters.

Now, with the recent introduction of densely packaged rack-optimized servers and blade servers, along with advances in software technology that make it easier to manage and exploit resources, it is possible to extend the benefits of clustered computing to individuals and small departments. With the Cluster 1350 (*Figure 4*), which is part of IBM Departmental Supercomputing Solutions, the barriers to deploying clustered servers—high price, complexity, extensive floor-space and power requirements—have been overcome. Clients with smaller budgets and staff, but with challenging problems to solve, can now leverage the same supercomputing technology used by large organizations and prestigious labs.

Figure 4. *Cluster 1350*

The Cluster 1350 is not a single product, but a flexible, integrated offering consisting of hardware and software delivered prepackaged in a rack. This approach can save you many hours of time compared with assembling, integrating, testing, tuning, managing and supporting the cluster yourself. It provides a validated configuration with a single point-of-contact for continuing support.

The Cluster 1350 combines the power of e326 and LS20 servers (as well as IBM @server xSeries® 336 (x336), x346, HS20 and JS20 servers) with sophisticated software, as well as storage and networking components, to create integrated, flexible offerings for HPC and commercial application environments. The Cluster 1350 includes cluster management software, such as the Linux Cluster Install Tool (LCIT), Cluster Systems Management (CSM) for Linux or AIX, xCSM for Linux, and General Parallel File System (GPFS). This software is what turns the collection of interconnected servers into a powerful yet manageable cluster for high performance and scalable computing. It helps you derive maximum performance and utility from the cluster, as well as ease systems management and enable rapid time-to-productivity.

Optional hardware, such as IBM TotalStorage® DS300, DS400 and DS4000 series storage servers and high-speed interconnects, including Gigabit Ethernet, InfiniBand™, Topspin and Myrinet, complete the package.

Cluster 1350 systems support up to a total of 1,024 compute nodes, including up to 64 storage nodes and a management node (with an optional second management node for redundancy), packaged in 11U, 25U or 42U racks.

IBM IntelliStation A Pro

Advantages:

- *Visualization workstation that doubles as a cluster/grid node*

- *Choice of task-tuned 3D graphics solutions from leading vendors*
- *Flexible, scalable workstation design*

High-performance computing consumes and produces tremendous quantities of data. The preferred means of interacting with data is graphically: visualizing the results of complex algorithms through sophisticated 3D graphics, revealing insights into the phenomena being studied. The graphical client of choice for HPC is the technical workstation, combining high-end computing performance with stunning graphics to rapidly conceptualize the results of computational work.

The IBM IntelliStation® Pro line is a feature-rich family of graphics workstations. The next-generation IBM IntelliStation A Pro (model 6217)—powered by single- and dual-core Opteron 200 Series 32-/64-bit processors (all dual-capable)—offers systems with up to four cores. A range of certified PCI Express (PCI-E) x16 graphics cards is available from NVIDIA and 3Dlabs[11], all with dual-monitor support, providing impressive graphics performance. Equipped with up to two processors, the scalable 6-slot x 6-bay IntelliStation A Pro is optimized for compute-intensive, graphics-rich applications. Like the e326 and LS20 servers, the A Pro architecture bypasses the primary performance bottleneck of other workstation designs—the front-side bus—providing up to 16GB of PC3200 DDR memory (8GB per processor) accessed at the *full processor clock rate* to manipulate massive datasets.

IntelliStation Pro workstations are developed in conjunction with xSeries servers. IBM workstation and server designs frequently share components, down to the motherboard level, resulting in server-class workstation reliability. This focus on reliability is reflected in other areas as well:

- Robust configurations of ECC memory
- RAID storage for customers who need higher data availability, reliability and security
- Diagnostic LEDs on select models that help you locate component problems in the event of an error
- IBM systems management tools, such as the IBM Director client included with every IntelliStation A Pro workstation

The IntelliStation A Pro is certified by leading 3D application vendors to visualize the universe—from the sub-atomic to the galactic. With the advent of dual-core Opteron processors, the A Pro has evolved into the premier platform for compute- and graphics-intensive applications, ranging from engineering to life sciences, petroleum exploration to digital content creation.

Additional Information

For more information on IBM @server and IntelliStation directions, products and services, visit our Web site at:

- http://**ibm.com**/servers/eserver/opteron for more on the Opteron processor-based **e326** and **BladeCenter LS20** servers
- http://**ibm.com**/servers/eserver/bladecenter for more on the **BladeCenter chassis** and options
- http://**ibm.com**/servers/eserver/clusters for information about @server **1350 clusters** using **e326/x336/x346** servers and **LS20/JS20/HS20** blade servers
- http://**ibm.com**/servers/intellistation/pro/apro for information about **IntelliStation A Pro** workstations

From any of the product pages, select **Literature** (in the blue navigation bar on the left) for links to the various types of documentation available.

[11] Currently offered adapters include: 3Dlabs Wildcat Realizm 800 with 640MB GDDR3 SDRAM (512MB unified, 128MB Direct Burst); NVIDIA Quadro FX 3400 with 256MB DDR3 SDRAM, 350 MHz GPU core; NVIDIA Quadro FX 1400 with 128MB DDR SDRAM, 350 MHz GPU core; NVIDIA Quadro NVS 280 with 64MB DDR SDRAM, 250MHz GPU core.

For more information about AMD Opteron processors, go to:

- *http://www.amd.com/opteron* for more on the AMD **Opteron processor**
- *http://www.amd.com/multicore f*or additional information about **multi-core AMD technology**

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

Produced in the USA
June 2005

Visit http://**ibm.com**/pc/safecomputing periodically for the latest information on safe and effective computing. Warranty Information: 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.