



32/64-Bit HPC and Cluster Computing with the Opteron-based IBM @server 325

*by Mark T. Chapman
IBM Server Group*

Executive Overview

Whether your organization is involved in petroleum exploration, computer-aided engineering, electronic design automation, the life sciences, or various other fields of scientific research, high-performance computing (HPC) isn't simply a necessity, it's the lifeblood of your organization. HPC customers face numerous challenges these days, and not just the cost pressures—the need to do more with less—that every IT shop is burdened with. More than that, there are those challenges that are more problematic in the HPC arena, which is always at the forefront of technology: the requirement for bleeding-edge performance, the prospect of 64-bit hardware architectures and their corresponding operating systems, the migration path to take users smoothly from 32-bit to 64-bit applications and, of course, the need to simplify the complexity of all this technology.

Sixty-four-bit servers have been around for many years, in the form of mainframes and UNIX®-based servers, such as the IBM @server™ pSeries™ and predecessors, yet in the Intel®-compatible world, 64-bit computing is a relatively new phenomenon, introduced with the Intel Itanium™ processor in 2000. The recent releases of the second-generation Itanium 2 and AMD™ Opteron™ processors have resulted in better price/performance measurements and a wider array of server offerings. This broader choice offers users greater flexibility, yet at the same time introduces some confusion into the marketplace. Suddenly it isn't a straightforward decision as to which 64-bit solution to buy. Should you stick with the tried-and-true UNIX solution with thousands of available 64-bit programs? Or are you better off going with an Intel-based solution, which may be much less expensive hardware-wise and provides Microsoft® Windows® compatibility, but that offers little off-the-shelf 64-bit software? Or, perhaps, you could save even more money with an Opteron based server, but at the risk of straying from the Intel fold.

IBM is the first major server vendor to offer an Opteron-based server, the IBM @server 325. This new addition means that IBM customers have three choices for 64-bit HPC computing (including our pSeries and xSeries® offerings). Each option has advantages and disadvantages relative to the others. Performing a complete analysis and comparison of the three product lines is beyond the scope of this short paper. Instead, its purpose is to briefly describe the e325 product and compare the business advantages of the server and the Opteron processor itself, and how they help address the customer challenges described above, with Itanium 2-based servers, such as the IBM @server xSeries 382 (x382) and 32-bit Intel Xeon™-based servers, including the x335.

In addition, a short set of decision maker's guidelines will help identify which choice is most appropriate for certain purposes. Regardless of which 64-bit solution you choose, you'll have the full backing of IBM—including our world-famous worldwide service and support infrastructure, our IBM Global Services consulting arm, and IBM Global Financing—to help you achieve your goals.

64-bit Technology

By itself, 64-bit technology is not a reason to jump in the air and click one's heels. It's what you as a customer can *do* with "64-bitness" that makes it interesting. For one thing, it removes the hard limitation of 4GB of memory access per application. Using 64-bit addressing, new servers can theoretically now directly access 16 exabytes (16 *billion* gigabytes!) of RAM. Of course, presently no Intel-compatible computer can physically hold that much, but whether it has 8GB or 16GB or 256GB, a single 64-bit application could access as much RAM as needed, given a 64-bit operating system. In addition, certain types of processing can be accomplished much faster using 64-bits at a time than 32-bit. But beyond these basics, there are specific architectural attributes of the 64-bit Opteron and Itanium 2 processors that affect performance.

64-bit Hardware Architectures

Like the Xeon and earlier Intel processors, Intel uses something called a front-side bus (FSB) to handle traffic between the processor and memory and I/O. Because the FSB operates at a much lower speed (400MHz or 533MHz, currently) than the processor clock rate (1.0-1.5GHz currently for the Itanium 2), there can be throughput bottlenecks when the FSB is congested with traffic. In contrast, Opteron replaces the need for the FSB with something called a HyperTransport™ Tunnel. This tunnel operates at the same clock speed as the processor itself, rather than being stepped down to a fraction of that speed, as the FSB in Intel processors is. This can translate into significantly higher bandwidth for the Opteron than for the Itanium 2.

Also, in an Opteron-based server memory has its own high-speed path to the processor, rather than having to share the FSB with other devices. Not only that, but 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 reads.) By comparison, adding a second Itanium 2 processor merely adds to the congestion of the existing FSB. This allows the performance of an Opteron-based server to increase linearly when additional processors are added. This is not the case with Itanium 2 servers, due in part to the increase in FSB congestion.

Another crucial architectural difference is that Intel teamed with HP to take the original Intel x86 architecture off on a tangent for 64-bit computing, replacing it with a new and entirely different instruction set. Opteron, on the other hand, extends the *original* Intel x86 core used in 32-bit Intel processors. What this means for users is that the Opteron can natively run both 32-bit and 64-bit software, concurrently, at extremely high levels of performance. Because Intel chose to optimize the Itanium 2 for 64-bit computing exclusively, the result is that while 64-bit performance is outstanding, 32-bit applications are forced to run in a special 32-bit emulation mode, resulting in subpar performance, even when compared with 32-bit processors. This disparity in performance tends to make the Itanium 2 unsuitable in an HPC environment that is still using 32-bit operating systems and applications.

64-bit Software

The normal life cycle for a server is to upgrade the hardware and software for as long as it makes economic sense, and then to retire the server and replace it with newer hardware. This is a fairly straightforward procedure when the new server is similar to the one it replaces: The 32-bit software you ran on the previous server will run as well, if not better, on the new server.

However, server upgrades can get bogged down in a quagmire of software versions when the upgrade isn't as straightforward. One day, when it's time to embrace 64-bit computing, you'll have to face the prospect of managing dual code bases: juggling 32-bit software on 32-bit servers and 64-bit software on 64-bit servers. How will you orchestrate the transition to go as smoothly as possible?

Using Intel-based servers, this migration may appear daunting enough to convince you to postpone the migration as long as possible. If you spend money on a number of shiny new Itanium 2-based servers you'll want to get the most from them right away. Yet, if you're still using 32-bit operating systems and applications, you won't receive the full value from your new 64-bit servers (especially given the Itanium 2's relatively poor performance when running 32-bit applications in 32-bit emulation mode).

As an alternative, Opteron is a 64-bit processor that runs 32-bit software with performance comparable to or, in some cases, *better than*, 32-bit processors. This means that you can buy an e325 today to run your 32-bit software, and then when you're ready for the switch to 64-bit computing, you can simply upgrade your operating system, applications and programming tools

to 64-bit versions and voila! You have 64-bit servers, without expensive hardware upgrades. 32-bit applications can be run on a 64-bit operating system until 64-bit replacements for the apps are available, and your budget allows, allowing for a staged migration. In many cases you may see a performance boost when using 32-bit applications and a 32-bit OS on an e325 (versus a 32-bit Xeon server), then another boost when you upgrade the e325 to a 64-bit OS running the same 32-bit applications, and then again when you upgrade the applications to 64-bit versions.

Performance Considerations

For some HPC customers, floating-point performance is all that matters; for others, integer results are paramount, and for still others, a balance in performance of both types of calculations is essential. Plus, of course, some types of computing lend themselves to being broken up into many small segments, appropriate for large clustered arrays, while other applications are more monolithic in nature and run best on one or a few large systems. Due to the differences in customer needs, a comparison between Intel Itanium 2-based servers, such as the x382 and x450, and the e325 is essential for determining the right type of server for you. With the rapid advances in processor technology, especially over the last few years, clock speed alone is no longer enough to identify which processor is fastest. Considerations like memory bandwidth and latency, interconnect capabilities and I/O throughput must also be taken into account, as well as whether the processor can run both 32-bit and 64-bit software effectively. After all, if the memory and I/O transfers are slow, the processor can't operate at peak efficiency; and if the processor can't run both 32-bit and 64-bit software well, how will you manage a smooth migration from one to the other?

In addition to raw performance benchmarks, price/performance is important to many customers. Comparing the performance of an e325 server, using the **Opteron 246** processor, to competitive servers using the latest versions of the **Itanium 2 (1.5GHz/6MB L3 cache)** and **Xeon (3.06GHz/1MB L2 cache)** reveals the outstanding price/performance of the e325. In the *SPECfp_rate2000* benchmark¹ of floating-point performance, the e325 (in a dual-processor configuration) posted a floating-point score that was **41% better** than that of the Xeon-based HP DL360 G3 server, and at a slightly *lower* cost. The same e325 score was **63%** of that posted by the Itanium 2-based HP rx2600 server, but at only **one-seventh the cost!**

In general, Opteron will outperform Itanium 2 with **32-bit** applications and for **64-bit Integer** use. Itanium 2 will outperform Opteron for **64-bit floating point** use, but at a far higher cost. Also, Opteron will handily outperform Xeon overall (for roughly the same cost), while offering the upgradability to 64-bit processing that Xeon lacks.

This paper doesn't have enough space to go into all the details of these processors, but in the *Decision Maker's Guidelines* section below, some of the major features of the Itanium 2, Opteron and Xeon are compared.

IT Budget Constraints

This technology is all well and good, but unless it fits into your budget structure, it can't help you. With the advent of the Opteron, there's an alternative to the Itanium 2, and one that can be much more cost-effective. If you buy AMD Opteron-based servers—such as the e325—for your 32-bit computing needs, you gain a lot of flexibility. Because Opteron is capable of running both 32-bit and 64-bit software natively (i.e., at full speed), it makes the transition to 64-bit computing easier, faster and less expensive for end-users and developers alike. You can use e325 servers to run your 32-bit software today, and then later upgrade them to 64-bit software without replacing the servers.

¹ HP rx2600 results: <http://www.spec.org/cpu2000/results/res2003q3/cpu2000-20030630-02322.html>; HP DL360 results: <http://www.spec.org/cpu2000/results/res2003q3/cpu2000-20030630-02353.html>; IBM @server 325 results: <http://www.spec.org/cpu2000/results/res2003q3/cpu2000-20030728-02418.html>.

e325 servers offer two other forms of investment protection: 1) They tend to cost significantly less for similar configurations than Itanium 2 servers², and 2) Should you find that your timetable for the migration to 64-bit computing has accelerated, you won't find yourself caught between server generations, where the current generation hasn't yet been depreciated and you aren't ready to spend the money to replace them. Simply upgrade the software earlier; your e325 hardware will be ready for the transition. In January 2003, a Giga Information Group paper³ showed that a staged migration of Opteron-based servers from 32-bit to 64-bit computing can, in many cases, save a company a considerable amount of money over a five year period: for both hardware and software conversion (including redeployment of some servers at the end of the normal life cycle).

Overview of the IBM @server 325

The e325 is designed to pack as much computing horsepower in as little space as possible. It's a space-saving **1U** server that ships with one Opteron processor, upgradable to two. The e325 was engineered specifically by IBM for the HPC market, with design points appropriate to that market:

- The mechanical design offers toolless rack-mounting and a toolless chassis, to simplify installation and servicing.
- The channeled airflow layout enhances system cooling.
- The e325 shares many options with xSeries servers, including SCSI hard disk drives, Converged Trays, and Advanced Connectivity Technology (ACT) cabling, to name a few.
- The e325 offers extensive standard and optional systems management features, including an integrated systems management controller, IBM Director, IBM Cluster Systems Management, Remote Supervisor Adapter II⁴, and others.
- The e325 is not overburdened with unnecessary features that complicate the design and add cost, such as an internal floppy drive. (There is a USB-attached drive available as an option.)
- The e325 is a supported platform in the IBM @server Cluster 1350 solution.

From a form factor standpoint, the **1U** height of the **2-way e325** provides twice the processor density of **2-way 2U** Itanium-based servers, such as the x382 and the HP rx2600. In addition, the compact **26"** depth of the e325 simplifies cable management and provides increased rack airflow vs. deeper servers.

Pricewise, for similar configurations the e325 is only slightly more expensive than the Xeon-based **x335**, and only a *fraction* of the cost of the Itanium 2-based **x382** or **x450**. This gives the e325 a large *price/performance* advantage over all three servers.

IBM Director 4.1x (planned for Q4/2003) software for advanced workgroup management is available for the e325 at no additional cost. IBM Director comes with a portfolio of tools for server management and increased availability. IBM Director provides a single uniform graphical interface for all systems management functions. The combination of IBM Director and the integrated systems management controller enables the customer to customize thresholds and monitor system components (for things like temperature, voltage regulation, etc.) to help maximize uptime.

Likewise, **IBM Cluster Systems Management (CSM)** can help to reduce the total cost of ownership for cluster environments by streamlining and simplifying the management of large numbers of servers from a single point of control. This is of particular value in large data centers or for

² Using public Web pricing as of August 8, 2003, the IBM e325 server with Opteron 246 processor cost **approximately 1/7** as much as a similarly configured HP rx2600 server with a 1.5GHz Itanium 2 processor.

³ "A Total Economic Impact Framework for Deploying AMD Processor-based Servers", January 2003, http://www.amd.com/us-en/assets/content_type/DownloadableAssets/TEI_for_AMD_Opteron_FINAL.PDF.

⁴ Support planned for Q4/2003.

customers who are running complex high-performance workloads, where managers are constantly challenged to meet ever-changing business demands and to improve service levels at the same time that IT budgets are shrinking. As customers evolve from using clusters for single-purpose workloads to employing cluster technology in grids or utility models, CSM can provide a path for growth and enhanced value.

Today, many businesses are building their own Linux clusters using commodity hardware, standard interconnects and networking technology, Open Source software and in-house or third-party applications. Many discover that this activity consumes considerable resources to assemble, integrate, test, manage and support the cluster. e325 customers can benefit from our extensive experience with clustered UNIX computers to minimize this complexity and risk. IBM designed the **Cluster 1350** to address these challenges. Using advanced IBM server nodes—including the e325, x335, x345, x360 and BladeCenter, proven cluster management software and optional high-speed interconnects, the Cluster 1350 brings together the best of IBM and ISV technology. As a result, the installation of a Linux cluster can be greatly accelerated and its management and support simplified.

Decision Maker’s Guidelines

HPC buyers are faced with the dilemma of whether to seek out computers that deliver the highest performance for their specific needs—even if it means going with a proprietary architecture; or to reduce capital outlay by using industry-standard computers, even if it means sacrificing top performance. Over the last decade, the shift has been increasingly *from* proprietary and specialized systems and *to* industry-standard systems which offer better price/performance. According to IDC⁵, in 1993 custom/proprietary processor-based servers accounted for **84.6%** of the overall HPC market, in terms of dollars spent. By 2002, custom/proprietary server spending had dwindled to only **5.9%**.

As with most IT decisions, the answer to the question of which is the best solution for a given customer is “it depends.” Every customer’s needs differ. This is why IBM offers a range of 32-bit and 64-bit products, each with its particular advantages over the others. Whether your needs are for maximum floating-point performance, large memory bandwidth, small form factor or PCI-X slots, we have the server that will best suit your needs.

To help you decide which solution is ideal for you, regarding both 32-bit and 64-bit rack-dense computing, *Table 1* compares some of the important features of the e325 with the x335, x382 and pSeries 615 offerings⁶, and offers suggestions as to which server is best for certain requirements:

Server Feature	@server Model	e325	x335	x382	p615
Processor type (2-way)		AMD Opteron	Intel Xeon	Intel Itanium 2	IBM POWER4+
Fastest clock rate (as of August 29, 2003)		2.0GHz	3.06GHz	1.5GHz	1.2GHz
Maximum HyperTransport Tunnel / FSB / system bus bandwidth		16GBps	4.26GBps	6.4GBps	12.8GBps
Maximum memory bandwidth (2-way)		10.6GBps	4.26GBps	6.4GBps	6.4GBps
Maximum I/O bandwidth (2-way)		12.8GBps	3.2GBps	6.4GBps	3.2GBps
Runs 32-bit Windows / Linux		Y/Y	Y/Y	N/N	N/Y
Runs 64-bit Windows / Linux / AIX® 5L		Y/Y/N	N/N/N	Y/Y/N	N/Y/Y
Maximum memory		12GB	8GB	16GB	16GB

⁵ “A Look at the HPC Server Market: A New Approach to the IDC HPC Balanced Rating”, IDC #29617, June, 2003.

⁶ Another 2-way IBM rack server not included here is the **2U 32-bit Xeon-based x345**.

Server Feature	@server Model	e325	x335	x382	p615
Memory type		PC2700 333MHz	PC2100 266MHz	PC2100 200MHz	IBM 200MHz
Number of full-length PCI-X slots		1.5	1.5	3	6
Number of 133MHz/100MHz PCI-X slots		0/1	1/0	1/0	4/0⁷
SCSI storage type		Ultra320	Ultra320	Ultra320	Ultra160
Maximum HDD storage		293.6GB	293.6GB	293.6GB	1.1TB⁸
Hot-swap power		No	No	Yes	Optional
Redundant power		No	No	Special bid	Optional
Hot-swap cooling		No	No	Yes	Optional
Redundant cooling		No	No	Yes	Yes
Form factor		1U	1U	2U	4U
Chassis depth		26"	25.7"	28"	20"
Diagnostic indicators		LEDs	Light path	LEDs	LEDs
C2T Interconnect™ cable chaining technology		No	Yes	No	No
Recommendations for:					
Lowest-cost 32-bit computing			✓		
Lowest-cost 64-bit computing		✓✓✓			
Best 32-bit integer performance (SPECint_rate2000)		✓			
Best 32-bit floating-point performance (SPECfp_rate2000)		✓			
Best 64-bit integer performance (SPECint_rate2000)		✓			
Best 64-bit floating-point performance (SPECfp_rate2000)				✓	
Best 64-bit price/performance		✓✓✓			
Highest availability / most redundancy				✓✓✓	✓✓✓
Highest rack density		✓	✓		
Largest memory requirements				✓✓✓	✓✓✓
Integration into IBM @server Cluster 1350		✓ ⁹	✓		
Integration into IBM @server Cluster 1600					✓

Table 1. Comparison of e325, x335, x382 and p615

Legend

✓ = Recommended

✓✓✓ = Significant advantage

⁷ The p615 has four times as many high-speed PCI-X slots as the e325, the x335 and the x382, but in a 4U cabinet, vs. 1U or 2U for the other servers.

⁸ The p615 has four times the internal storage capacity of the e325, the x335 and the x382, but in a 4U cabinet, vs. 1U or 2U for the other servers.

⁹ Available Q4/2003.

As you can see, although the e325 offers significant advantages in terms of performance, price performance and rack density, there are some situations where the x382 would be a better fit, such as where huge memory resources are required, the highest possible availability is essential, or where the utmost in floating-point performance outweighs the higher cost factor for the x382.

If your most pressing requirement is price, and you don't expect to have significant 64-bit processing needs for years to come, then the x335 would be the best choice of the three, allowing you to defer the 64-bit decision until later.

The Complete Solution

e325 servers have the ability to run both 32-bit and 64-bit software effectively, both individually and concurrently, which offers a lot of flexibility to customers. But there is much more to simplifying your HPC/clustering environment than that. There are bigger issues, such as:

- Can your solution vendor provide you with expert preparation and planning tools and expertise?
- Does your solution vendor have relationships with leading HPC software vendors?
- Can your solution vendor help you manage a multisystem/multilocation installation?
- Does your solution vendor offer effective provisioning and deployment tools to simplify the task of rolling out an organization-wide platform?
- Does your solution vendor test and support a high-speed interconnect method (for example, Myrinet)?
- Can your solution vendor deliver the complete HPC solution you need (for example the IBM @server Cluster 1350)?

IBM has been a leader in the clustering and deep computing world for years. We have technologies that scale up through midrange and high-end servers, and which increasingly are migrating to our industry-standard servers. By lowering the cost of computing without compromising quality and service we provide you with tremendous value. Increasing your productivity to new levels will position your business for efficient future growth.

IBM has long experience providing complete solutions for complex system environments and has one of the largest networks of experts skilled in capacity planning, multiplatform integration, system and application development, systems assurance and performance tuning. We can help you assess your system options and implement the best @server solution for your environment.

Summary

The advent of servers such as the e325 that use the high-performance, low-cost AMD Opteron processor, has opened 64-bit computing to all customers, not just those with deep pockets. The e325 server's ability to run 32-bit software today with leading-edge performance and then to easily convert to running 64-bit software when you are ready to do so eliminates the need for costly hardware upgrades and minimizes the headaches of supporting dual code bases for months or years during the transition.

If your needs require an Itanium 2-based server, such as the x382; or your budget limits you to a lower-cost high-performance 32-bit server, such as the Xeon-based x335; or you prefer the p615, which supports the industry-strength UNIX[®] implementation of IBM AIX 5L; you are still covered by the broadest line of 32-bit and 64-bit servers in the industry: IBM @server. And every @server system is backed by IBM's famed service and support, consulting services and financing. You can't lose no matter which course you steer through treacherous IT waters.

Additional Information

Visit <http://ibm.com/pc/us/eserver/opteron> (or call **1-888-SHOIBM**) for information about IBM @server e325 products and services, including part numbers and prices for servers, racks, storage units and other options.



© IBM Corporation 2003

IBM Server Group
Department U2SA
Research Triangle Park NC 27709

Produced in the USA
8-03
All rights reserved

IBM, the IBM logo, the e-business logo, Active PCI, Active PCI-X, AIX, BladeCenter, C2T Interconnect, Chipkill, NetBAY42, Netfinity, Predictive Failure Analysis, pSeries, X-Architecture, XceL4, XpandOnDemand, Wake on LAN, WebSphere and xSeries are trademarks of IBM Corporation in the United States and/or other countries.

AMD, HyperTransport and Opteron are trademarks of Advanced Micro Devices, Inc.

Intel and Pentium are registered trademarks and Itanium and Xeon are trademarks of Intel Corporation.

Linux is a registered trademark of Linus Torvalds.

Lotus and Notes are trademarks or registered trademarks of Lotus Development Corporation and/or IBM Corporation.

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

UNIX is a trademark in the United States and other countries licensed exclusively through The Open Group.

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

All offers subject to availability. IBM reserves the right to alter product offerings and specifications at any time without notice. IBM is not responsible for photographic or typographic errors. 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 document contains information relating to future or anticipated releases of products and represents IBM's current intentions, goals and objectives. The information in this document is subject to change or withdrawal without additional or prior notice.

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.

IBM @server xSeries servers are assembled in the U.S., Great Britain, Japan, Australia and Brazil and are composed of U.S. and non-U.S. parts.

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.