

PUND-IT RESEARCH

Marketplace Update

August 1, 2006

Business Performance Computing – Changing the Commercial Equation for the x86 Industry

By Charles King, Principal Analyst, Pund-IT, Inc.

Pund-IT, Inc.
2776 Sulphur Drive
Hayward, CA
U.S.A. 94541

Phone: 510-909-0750
Fax: 510-886-4937
charles@pund-it.com
www.pund-it.com

Overview

Over the past few years, mainstream businesses have begun to adopt high performance computing (HPC) technologies to deliver cost-effective solutions for a widening range of commercial applications. IBM, a long-standing leader in HPC and supercomputing innovation, has enjoyed particular success in early adopter commercial HPC markets including Monte Carlo applications, Computer Aided Design (CAD), and simulation/testing solutions. This evolution of HPC for more mainstream applications is defining a segment called Business Performance Computing (BPC).

AMD is another vendor that has inspired significant shifts in the HPC marketplace, largely as a result of its Opteron processor. Opteron was initially lauded for its unique (for the time) support of 64-bit extensions, an AMD effort that fundamentally altered the market for x86-based servers. But the notable price/performance and throughput of the Opteron architecture also helped speed the evolution of HPC. Recognizing these singular qualities, IBM became the first major vendor to add Opteron based servers (the e325) to its HPC portfolio.

Commercial BPC solutions can be particularly valuable because they offer easy scaling and cost-efficient upgrades. In addition, the price structure of x86 solutions lends itself to typical business IT models, as do the availability of standard operating systems and ISV offerings. Businesses, however, require not just pure, cutting-edge technical performance, but solutions that deliver the reliability and availability necessary to support everyday business applications and processes.

As the demand for BPC grows, opportunities are developing in mainstream business applications and in emerging markets including product validation, business intelligence, pattern recognition, IPTV, and video on demand (VoD). IBM is leading the charge in this market by focusing on new servers that blend the performance of the company's well-established, innovative availability and manageability features with the performance of AMD's Opteron processors. As a result, IBM's new System x and BladeCenter solutions deliver BPC solutions with the price and ease of use computing customers enjoy from commercial solutions and the performance they need to existing and emerging critical business applications.

Business Performance Computing – Changing the Commercial Equation for the x86 Industry

By Charles King, Pund-IT, Inc.

Information technology (IT) is evolutionary by nature, driven by competitive forces within both vendor and developer communities, and by demand in the greater marketplace. But while many IT changes inspire little more than yawns from jaded industry watchers, some shifts are notable enough to jerk even the most somnolent awake. Developments in the supercomputing and high performance computing (HPC) arena are prime examples of this. For decades, supercomputers were typically sizeable monolithic systems of the sort that served as inspirations for HAL in “2001 – A Space Odyssey,” but contemporary HPC installations are significantly different creatures than that morose beast.

Highly flexible and sophisticated clustering applications sparked a new generation of supercomputing and HPC solutions that have, quite literally, taken the market by storm during the past half decade. How radical is this change? In June 2001, clusters represented a mere 6.4% of the systems on the top500.org list of the world’s highest performing supercomputers. In the June 2006 top500.org announcement, 364 clusters constituted 72.8% of systems on the list. Equally important however, is the affinity between high performance clustering solutions and the x86 processor architecture. In June 2001, the top500.org list included just eleven IA-32-based systems. Today, various x86 flavors including Intel IA-32 and EM64T, and AMD’s Opteron provide the brains and horsepower for 345 (or 95%) of the clustered systems on the Top500.org list.

It would be a mistake to suggest that Top500.org represents supercomputing and HPC as a whole, but it is reasonable to assume that trends driving the higher end of the market also influence the lower end, especially those related to accessible and affordable technologies such as commercial-off-the-shelf (COTS) x86 servers. In addition, other factors have a profound influence on x86 HPC systems, particularly the use of Linux applications. This is hardly surprising, given the key role commodity x86 servers played in the early development and uptake of Linux business applications. Linux is also dominant force in academic IT environments and labs, where much of the more esoteric supercomputing and HPC work is being performed. At the same time, Microsoft’s growing interest in HPC is likely to be a significant driver of these solutions among smaller and mid-market organizations.

Unreal HPC Performance Meets Real World IT

Good enough, but do these issues have broader implications? Absolutely. Maturing scale out x86 HPC environments are changing the way companies do business in areas ranging from automobile design and crash testing to Monte Carlo simulations that offer insight across a range of risk analysis applications. In addition, the price/performance of x86 scale out solutions continues to drive software evolution into new commercial application areas in a trend known as Business Performance Computing (BPC). As these technologies continue to mature we expect them to make further inroads among enterprise customers and in broader markets and commercial business applications.

IBM has long been a leader in supercomputing and HPC, and the company consults on or spearheads numerous high-end research and academic projects. IBM-based systems make up just under half of the current Top500 list. But those numbers do not tell the entire story; as a whole those 239 IBM-based systems deliver nearly three times more total supercomputing performance than second place HP’s 157 systems. Along with these heavy lifting high-end systems, IBM is also at the forefront of commercial BPC in developing industry-focused solutions with a variety of developer partners including AMD.

AMD is one vendor that has developed notable traction in performance computing. The obvious key to this success has been the company's Opteron processor, which fundamentally changed the face of x86 IT solutions and established AMD as an innovative leader in this space. In June 2001, top500.org ranked only one AMD system but the most recent list boasts 80 AMD Opteron installations two of which are in the top 10, four in the top 20, and 21 in the top 100. IBM Opteron installations account for 31 of AMD's current top500.org crop, and 13 of those systems reside in industrial/commercial settings.

IBM and AMD – A Dynamic Partnership

Top500.org illustrates some obvious results of IBM and AMD's common efforts, but the companies' partnership is considerably broader and deeper than many realize. In 2003, IBM became the first Tier 1 vendor to formally embrace Opteron, and the company's eServer 325 solutions set the tone for what customers and the market came to expect from AMD.

But IBM's plans for Opteron did not begin and end with the e325. Since then, the company has continued to leverage the processor's notable price/performance and technical leadership in solutions including the updated e326m, the industry's first Opteron dual core blade server (the AMD Opteron LS20 for IBM BladeCenter), the IntelliStation A Pro for high end graphics and visualization, and the Linux-focused Cluster e1350. Further, in August of this year, IBM will be the first to announce a new System x portfolio based on the next generation of the Opteron processor, with five innovative new offerings that include unique, industry leading solutions for both HPC and BPC.

IBM and AMD collaborate across a wide range of commercial and industry efforts that leverage the x86 architecture value chain. The companies have committed to numerous R&D, design, and process technology agreements, including their joint development alliance focusing on 65 nanometer chip development efforts at AMD's fabrication facilities in Dresden, Germany, and the IBM Semiconductor Research and Development Center in East Fishkill, N.Y. Both companies also mutually support industry and standards efforts including Blade.org, the HyperTransport Consortium, and the Green Grid.

Enhancing the Business Value of HPC

Does the IBM/AMD relationship reflect any larger trends within the HPC and BPC sectors? We believe so. While there are notable differences between HPC and BPC, there are also significant similarities. At one level, HPC aficionados demonstrate classic early adopter behavior while BPC reflects early majority trends among commercial interests. In short, HPC tends to thrive in the technically sophisticated realms of research and university labs such as Lawrence Livermore, Los Alamos, Cal-Tech, and the University of Washington. BPC adherents, on the other hand, are searching for a technical and competitive leg up in business application and solution performance.

Both groups are on the lookout for price/performance benefits, but for different reasons. Labs and universities deal constantly with critical budget issues, but HPC-driven projects are often short term or limited in focus. In such circumstances, the flexibility of these installations allows researchers to easily reconfigure them for new projects. In addition, labs and universities appreciate solutions that support cost-effective CPU- and I/O-intensive performance, large memory scalability, and quick memory access.

For all their similarities, HPC and BPC users tend to bring notably different IT skill sets to the table. With academic training and backing from some of the world's great universities and labs, HPC-focused researchers typically enjoy strong technical expertise including the ability to configure systems and write or customize software as needed. Many businesses have limited technical expertise, particularly in BPC, which necessitates the hiring of specialists and consultants, thus increasing overall costs.

While businesses want the enhanced performance provided by BPC solutions, most also need cost-effective products that support easy to use, dependably replicable solutions that fit their in-house skill sets. Bottom line? For commercial clients to get the most out of BPC, it is critical that solutions incorporate established system management and RAS features such as those provided by IBM.

BPC – Potential and Emerging Opportunities

Beyond the CAD, simulations, and Monte Carlo applications where these technologies have made initial inroads, what other commercial areas offer opportunities for BPC solutions?

- **Business Intelligence** – Leveraging information for maximum business value is a prime directive for enterprise IT, but the job is getting harder and harder as the sheer volume of data and information sources continues to rise. The ultimate goal: How about real time business intelligence (BI) that delivers the information managers need for accurate predictive decision making, and business process management (BPM) solutions that ensure organizations avoid missteps? BPC solutions offer the muscle needed to make such BI solutions real.
- **Product Design and Manufacturing** – One of the classic areas for demonstrating the real world IT innovation is manufacturing, where technology is delivering tangible value across a wide range of business processes. A growing number of manufacturers are leveraging HPC systems for designing and testing products from simple snack foods to complex machinery to efficient, cost-effective packaging. Clustered BPC solutions can supply the players in this space the tools they need to compete and succeed.
- **Risk Assessment** – Since Hurricane Katrina, insurers and others have become increasingly sensitive to the role accurate risk assessment plays in their businesses. While such solutions have long been common among larger insurers, commodity-based BPC installations can support powerful, affordable risk modeling tools for the organizations that need them. With the incidence of extreme weather events expected to rise, so will the demand for such solutions.
- **IPTV/Video on Demand** – Media qualifies as the Wild West of IT, jumbling together well-established entertainment barons and cable and satellite entrepreneurs developing next generation Internet-enabled television programming (IPTV) and Video on Demand (VoD) services. In addition, telecoms also want to grow their capacity within the same footprint, or central office (CO), to increase ROI. For such businesses, BPC can offer a serious and entirely necessary competitive advantage. The price/performance of these solutions makes them attractive to both smaller players looking for an edge in the game and to larger organizations that need to contain costs amidst cutthroat competition.
- **Pattern Recognition** – Once relegated to the realm of science fiction, technologies such as facial recognition tools are becoming increasingly common for addressing security concerns ranging from anti-terrorism efforts to curtailing shoplifters. But effective solutions requires IT horsepower, making these solutions prohibitively expensive for many public and private sector organizations. COTS servers can effectively lower the bar and make these solutions available and affordable to greater numbers of users, while providing the memory-intensive performance they require.

IBM - Mixing BPC Leadership and Innovation

AMD's leadership in 64-bit x86 technologies makes the company a solid BPC partner, and Opteron's inherent architectural advantages make it an obvious choice for these solutions. As a result, Opteron-based systems offer businesses the means to take quick advantage of existing and emerging BPC opportunities.

But while Opteron supplies the core of these solutions, not all servers are created equal. IBM offers numerous X-Architecture and other innovations in its new System x and BladeCenter servers that define the literal difference between good and great Opteron solutions. Among these features:

- **Innovative Linear Scalability** – IBM's clever use of HyperTransport allows customer to easily scale its new LS41 from dual socket to four socket, and its new x3755 system offers memory and CPU scalability with the flexibility of 1, 2, 3 and 4 socket configurations.

Xcelerated Memory Technology – This IBM technology uniquely helps customers optimize system performance via large, fast (667 MHz) memory subsystem. As a result, IBM's x3755 and the x3455 systems (which support 128GB and 48GB of internal memory respectively) help customers enjoy performance advantages that competing solutions do not and can not match.

IBM eXtended I/O – For clients looking for maximum flexibility and performance in I/O connectivity and systems management, IBM's eXtended I/O supports HTx, PCI Express, and PCI-X. Given the fluidity of the fast-changing I/O market, this innovative approach also provides investment protection for IBM's System x and BladeCenter customers.

Cool Blue and PowerExecutive – Long an innovator in power and cooling solutions, IBM's Cool Blue strategy helps customers take a longer view and better advantage of these constantly evolving technologies. The company's new PowerExecutive solution allows customers to monitor actual power consumption from individual servers to an entire datacenter, providing the means for intelligently managing IT resources and maximizing datacenter performance. Next on the Cool Blue agenda: Integrating power and virtualization management.

IBM System Cluster e1350 – Scale-out clustering solutions offer a widening variety of business benefits, but many companies lack the time and skills to deploy them. IBM's ready to run Cluster e1350 is designed to take the complexity out of cluster installation by delivering highly scalable solutions for HPC, commercial applications, and server consolidation that provide a single point-of-control for simplified management and enhanced availability.

IBM's New Opteron-Based Solutions

IBM's new System x and BladeCenter offerings provide customers a mix of entirely new and thoroughly updated BPC and HPC solutions:

IBM System x3455

The company says the new System x3455 is the fastest, leanest compute node on the planet. That may be a bit of an overstatement, but the x3455's combination of low cost, dual socket Opteron processors, up to 48 GB of DDR memory, eXtended I/O features, and Xcelerated Memory Technology (all packed into a 1U box) make it an excellent choice for the high computational workloads and memory-intensive applications common to HPC clusters in areas including digital rendering, electronic design, and scientific and technical computing. The x3455 is available preconfigured and pre-tested for Linux Cluster e1350 installations, and is also a cost-effective solution for small business Microsoft clusters and database applications.

IBM System x3655

IBM describes the new x3655 as a Business Performance Server that mixes the reliability of the company's mainframe solutions with dual-socket Opteron application performance. The x3655 offers up to 64GB of DDR2 memory, internal disk and tape storage options, eXtended I/O, and IBM Power Executive and Predictive Failure Analysis features. Additionally, the x3655 provides users a range of IBM availability features including redundant power supplies and fans, standard RAID, and systems management tools. Given its capacious memory and I/O flexibility, the x3655 is an especially appropriate choice for critical databases, Business Intelligence applications, web serving and for other demanding areas including IPTV, VoD, and security solutions.

IBM System x3755

The new 4U x3755 raises the bar considerably for large memory HPC compute nodes. In fact, IBM focused its x3755 efforts on establishing leadership in Performance/Watt, Price/Performance, and I/O performance. The result is a rack-optimized solution that offers faster memory access, superior memory performance, and more flexible I/O than competing Opteron products. The x3755 offers performance technologies including four socket Opteron processors, up to 128GB of DDR2 memory, SAS internal HD storage, configurable I/O slots, HTx I/O support, notable 3-way scaling, and Xcelerated memory Technology for exceptional memory throughput. In addition, the system deliv-

ers enhanced availability via IBM features like redundant power supplies and fans, standard RAID, and system management solutions. Overall, the x3755 is a robust HPC package ideal for scientific and technical computing environments, and demanding applications including computer aided engineering (CAE), crash simulations, financial analysis, and weather simulations.

BladeCenter LS21 and LS41

IBM's latest Opteron-based blades offer something updated and something new. The LS21 is the follow-on to the company's LS20, the first Opteron-based blade, and offers a lean, flexible solution that leverages the BladeCenter's centralized management and power efficiency capabilities. The LS21's dual socket Opteron processors and 32GB of DDR2 memory make it a good choice for highly scalable scientific and technical computing environments, and for front and mid-tier business applications including financial services, security, and SQL databases. The LS41 is a beefy solution with 2/4 socket Opteron processor options, 64 GB of DDR2 memory, and two internal SAS hard drives and RAID, and is optimized for scalable enterprise workloads and large databases. As a result, the LS41 is a fine fit for HPC clusters and for business solutions including data marts, data warehouses, ERP, and SQL databases.

Real World BPC – Customer Profiles

Red Bull Racing

Red Bull Racing is one of two Formula One racing teams (the other is Scuderia Toro Rosso) owned and operated by the Red Bull beverage company. Like all Formula One teams, Red Bull Racing is on the lookout for any competitive edge, and the team has found one in IBM's System x solutions.

Red Bull Racing had a long interest in HPC, and initially created a crude cluster of graphics workstations to bolster design efforts. But in 2004, the team realized that it needed to make a substantial HPC investment to provide computational fluid dynamics (CFD) enhancements for its two wind tunnel groups' aerodynamic simulation efforts.

After examining solutions including conventional Cray systems and Itanium clusters, Red Bull Racing chose a 128 core IBM e325 cluster. According to Nathan Sykes, Red Bull's CFD Manager, the IBM/Opteron solution offered substantial cost/risk/return benefits compared to other offerings, as well as hands-on service from IBM and AMD.

Since then, Red Bull has expanded its reliance on IBM, upgrading to an e326 with Infiniband cluster that leverages 256 Opteron cores on 128 sockets. Sykes noted that IBM eXtended I/O includes HTx technologies that provide highly flexible support of the Infiniband interconnects the team prefers, helping to achieve great I/O throughput. This fall, Red Bull plans to add a 512 core x3455 cluster to its CFD arsenal.

How has this investment paid off? According to Sykes, Red Bull's HPC environment essentially complements the wind tunnel teams' work by helping to enhance results, pinpoint problems, and speed new developments. In addition, aerodynamic efficiency and stability are areas that affect virtually every other Formula One system, from engine and fuel efficiency to brakes and safety design.

Along with delivering the competitive advantages Red Bull Racing needs, the IBM system performs with notable reliability; in three and a half years of operation, there have only been two minor system failures. Sykes noted that design efforts in the racing world tend to presage those in greater automobile industry by two to three years. Even as Red Bull's own CFD efforts expand, he expects clustered HPC solutions to play an ever-increasing role in optimizing the design of commercial vehicles.

Move, Inc.

Founded in 1996, Move, Inc. (formerly Homestore, Inc.) describes itself as the owner of the No. 1 network of home- and real estate-related Web sites. The company's resources include the Realtor.com, Move.com, and Moving.com sites; Home Plans home design products, the Top Producer contact management system for real estate agents, and Welcome Wagon, the 78 year old marketing company for new home owners.

Move, Inc uses performance computing to support and enhance a variety of business results. The company is extremely Web-centric, utilizing IT solutions across both Internet-focused Web app and database resources and all its business and back office processes. The company's interest in IBM's System x solutions arose during planning of a significant shift in the company's IT resources; the relocation of its entire datacenter (including 2000 Dell rack servers) from Thousand Oaks, CA to Phoenix, AZ. Move, Inc. had chosen to site its new datacenter in an existing facility, but reducing power consumption was crucial to implementing the project.

Vincent Stephens, Move's VP of Technology Operations, said that the company decided on to shift much of its operations to a blade server environment to reduce its server footprint and help conserve energy. After some consideration, Move began head-to-head testing (mostly focused on SQL performance) between IBM's Opteron-based LS20 blades and Dell 1855E blade solutions. According to Stephens, the IBM blades "blew Dell away" in both system performance and power efficiency—what business performance computing customers require.

Currently, Move, Inc. is running 16 BladeCenter chassis fully loaded with LS20 servers and has prepared its facility to add six more chassis soon. During the relocation process, the company also reduced its IT environment from 2000 to fewer than 600 total servers, largely as a result of physically and virtually (VMware) consolidating applications and workloads. Move plans future consolidations via additional deployments of VMware's ESX Server. The company also plans to launch additional Web sites and to add further applications as business processes require.

Along with increased system performance and power efficiency, the company has enjoyed a wide range of other benefits from choosing the AMD Opteron LS20 for IBM BladeCenter. Move's Web-based business is extremely time-to-market sensitive. Stephens says that using former datacenter processes, launching the new Move.com site (which debuted in May 2006) would have taken 15 to 18 months of effort instead of the 4-5 months it actually required. Even simple processes are easier; typical server deployment time has shrunk from five days to two hours.

According to Stephens, after factoring in cooling and energy requirements, increased datacenter efficiencies, and the superb support Move, Inc. receives from IBM and VMware, choosing IBM's Opteron-based LS20 blades was not just a good technical decision, but also a great business decision.

Mission Accomplished?

In 2003, IBM's e325 servers inspired a number of significant shifts in the IT marketplace, including the evolution of clustered high performance computing solutions. In becoming the first major vendor to adopt Opteron, IBM demonstrated its continuing leadership in innovative HPC and super-computing development. The company's focus on Opteron also demonstrates the deeper partnership between IBM and AMD, reflected in a variety of continuing, mutually beneficial company and industry efforts.

Opteron also provided the means for IBM and other vendors to deliver cost-effective Business Performance Computing (BPC) solutions for a range of commercial applications including Monte Carlo applications, Computer Aided Design (CAD), and simulation/testing solutions. As the demand for BPC expands, we expect further opportunities to develop in mainstream areas such as business intelligence, manufacturing, risk assessment, IPTV and Video on Demand, and pattern recognition. But while commercial BPC solutions can be particularly valuable for many companies, organizations

also need solutions that deliver the reliability and availability necessary to support everyday business applications and processes.

Not all BPC solutions are created equal, and IBM technologies including Xcelerated Memory Technology, eXtended I/O, and Cool Blue features such as PowerExecutive define the difference between good and great Opteron-based servers. But the company's new products also illustrate more than the depth of IBM's technical innovation and the breadth of its relationship with AMD. Most importantly, the new System x and BladeCenter offerings demonstrate IBM's notable expertise in creating Business Performance Computing solutions that deliver the price, performance, reliability, and ease of use demanded by organizations of every kind.

© 2006 Pund-IT, Inc. All rights reserved.

About Pund-IT, Inc.

Insightful industry analysis and balanced guidance have never been more important for IT vendors and their partners and customers. Pund-IT's emphasis is on understanding product and technology evolution and interpreting the effects these changes will have on business customers and the greater IT marketplace. This report is the result of sponsored research developed by Pund-IT, Inc., which believes its findings are objective and represent the best analysis available at the time of publication.