THE CLIPPER GROUP
Navigator™

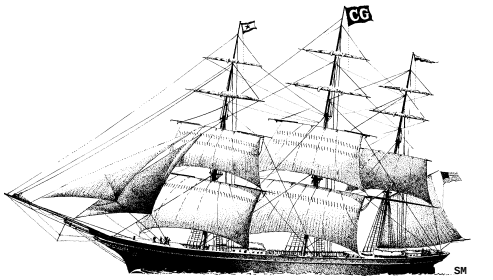

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2006073 August 22, 2006

IBM Makes Big Blue "Green" — Using Open Systems to Reduce Data Center TCO

Analyst: David Reine

As the cost of energy continues to rise, we try to become more innovative in the ways that we conserve energy, but in the words of *Kermit the Frog*: "It's not easy being green". Truer words could not be spoken, as we try to make our lives more "green". We try to innovate as much as we can. We implement energy conservation ideas in our homes, from simple things like turning off the lights when we leave a room, to replacing inefficient windows and installing wall insulation to save heat and A/C, to installing a solar heating system. We improve the efficiency of our transportation by trading-in our gas guzzling SUVs for hybrids that are more efficient. We even "go back to the future" and carpool to work or take public transportation as we did 30 years ago when Kermit was a mere tadpole. Today, Kermit has proven that being green can be "cool".

Conserving energy cannot stop in our homes and automobiles, however. Every enterprise must implement energy conservation policies in the workplace, from the power plant that heats and cools the buildings to the data center that drives the information technology engine. Today's data center employs a myriad assortment of inefficient server appliances throughout the enterprise - servers that predate our current energy crisis. The enterprise must improve the performance and utilization of these servers while, at the same time, implement the technological innovations that are being developed to reduce the operational costs of the data center and reduce total cost of ownership (TCO). This innovation is now available in the high performance CPU that is at the heart of almost every PC and server in the data center. Innovations in the area of energy consumption, performance, and virtualization have lead to major improvements in the utilization of server resources. Following the lead of chip technology innovators such as IBM and Sun, AMD introduced the first dual-core x86 processor in 2005, *Opteron*. In order to maintain its lead in the microprocessor industry, Intel followed with the introduction of a dual-core *Xeon*. Advances in dual-core technology continue, but the processor is only one component in a complex server solution. That requires innovation in memory and I/O architecture, and in cooling techniques that can dissipate the heat generated before it becomes a problem.

One company that has tried to stay ahead of the energy crisis in the data center is IBM. With almost a half-century of mainframe data center experience, IBM has adapted the lessons learned to creating energy-efficient open system data centers. Building upon a ten-year partnership with AMD, including three years with Opteron, IBM has now expanded its *System x* product line with five new Opteron servers, both rack and blade, and a new energy-efficient server-surround strategy - called *CoolBlue* - to reduce data center TCO. To see how Big Blue can use CoolBlue to save you real $green$, please read on.

IN THIS ISSUE

Open System Data Center Issues

Today's open system data center consists of a vast number of inefficient, dual- and quad-processor servers, often clustered together in a scale-out architecture, based upon single-core x86 technology from Intel and AMD. Combined with standard, open operating environments such as *Windows* and *Linux*, open systems servers provide both a lower cost structure and the interoperability that the enterprise demands. Unfortunately, these servers consume an excessive amount of energy due to poor utilization of CPU resources - powered on but unused - and poor power management. The typical data center server utilizes less then 20% of its available compute capability. **This has to change!**

Enterprises continue to place additional burdens on their IT infrastructure in order to compete more effectively with an ever-expanding application set. These new high-performance computing (HPC) solutions include numerical and memory intensive applications applied to everyday business functions, requiring higher levels of computational processing power and availability than before. Combined with system management features, they support requirements for business intelligence, enterprise resource planning (ERP), digital media, and modeling and simulation. The data center needs to avail itself of the latest technology to be able to install servers with multi-core technology, taking advantage of increased compute capability, faster memory and a virtualized environment, enabling multiple concurrent sessions running a mix of Windows and Linux applications. This will allow the data center to utilize significantly more of its installed resources. It will also improve the return on investment and lower the total cost of IT ownership (TCO).

Older, single-core servers simply run too hot. The energy required to power and cool the IT infrastructure, and cool the data center, itself, represents 25% to 40% of a data center's operating budget. Consolidating mission-critical and business-critical applications on fewer, more powerful, multi-core servers will enable the enterprise to implement virtualization techniques that are now readily available and will permit the data center to reduce administrative overhead by reducing the number of servers required. However, simply replacing a single-core CPU with a multi-core processor will not solve the TCO problem by itself. Additional architectural changes must be made to contain the demand for energy while providing additional computational capability effectively. Whichever vendor can best combine the performance advantages inherent in multi-core processors with architectural innovation in order to control energy usage and heat emission will command center stage in the x86 arena. IBM has been doing its best to fill that role.

IBM and AMD Partnership

In 1997, IBM recognized the need for a new approach to system architecture and formed an alliance with AMD for a new processor technology. That alliance has matured over the past decade, applying the AMD approach to the server environment, utilizing AMD's expertise in x86 microprocessor design and IBM's understanding for a total server solution.

In 1999, IBM and AMD launched discussions that would result in a development agreement in 2002, and the announcement of the Opteron microprocessor in June 2003. This was a true partnership with literally hundreds of engineers from both sides co-located, working together to improve the process technology. (IBM has even adapted these advances to the *POWER* architecture to improve *System p*.) This partnership resulted in the availability of the *eServer 325* in November of that year, and availability of dual-core Opteron servers in 2005[1]. In the summer of 2005, IBM shipped additional rack models, with its first dual-core Opteron blade, the *LS20*, which shipped in July. Now, IBM has extended its investment with AMD, introducing a new family of *System x* rack-mounted servers and *BladeCenter* blades, based on *Opteron Revision F*. This marks the marriage of AMD performance with IBM innovation: a new data center technology to limit and manage the amount of heat being disbursed into the data center. The architecture that IBM incorporated into System x servers is called *CoolBlue*.

[1] See **The Clipper Group Navigator** dated May 27, 2005, entitled *Scale-Up and Scale-Out Architectures – IBM Provides Choice with the xSeries*, available at http://www.clipper.com/research/TCG2005031.pdf

System x and BladeCenter Additions

IBM has announced a complete new line of dual-core Opteron Rev. F servers: three new rack-mounted models and two new blade servers for BladeCenter (See Exhibit 1, below.) These new servers represent the latest advancements in high-performance computing and provide the enterprise with more choice for x86 processing, complementing the line of Intel *Xeon* servers already available for System x. In addition, all Rev. F servers are designed to be upgradeable to four-core Opteron processors, when they become available.

Revision F incorporates on-chip virtualization, an updated memory controller that addresses DDR2[2] memory, and support for an HTX port, permitting the addition of add-in cards containing co-processors for a variety of tasks. These could include heavy-duty processing of XML files or the implementation of faster network connections. These new Opteron-based systems deliver 21% greater performance than the previous models, within the same power envelope, utilizing new power supplies that are 90% efficient. Near linear scalability enables 30% greater performance per watt.

Exhibit 1 – New System x Models

Rack Mounted

- System x3455 – A 1U, 2-socket HPC node ideally suited for HPC clusters engaged in database processing and scientific computing, scaling to 48GB of memory;
- System x3655 – A 2U, 2-socket Business Performance Server built for demanding I/O applications such as IPTV and Video-on-Demand;
- System x 3755 –A 4U, 4-socket HPC compute node with up to 128GB of memory and outstanding price/performance and scalability for large enterprises with demanding applications;

BladeCenter

- LS21 – A 2-socket HPC Blade optimized for performance computing, such as in financial services;
- LS41 – A scalable 4-socket enterprise blade with innovative "snap-in" design for on-demand growth in HPC clusters.

Source: IBM

IBM is not the only Tier 1 vendor to offer scalable x86 servers, however. There are others, such as HP and Sun, which have Opteron-based servers available. Opteron is a commodity architecture and AMD will provide its technology to all who ask. In order to distinguish one Opteron-based server from another, you need more than a fancy logo. IBM has optimized these models for *business-performance computing*, leveraging high-performance computing capabilities to address mainstream business needs using industry-standard technology. IBM has incorporated AMD's Opteron within its own *X-Architecture* with mainframe-inspired capabilities, superior memory innovation, and I/O leadership. They have complemented AMD's *Direct Connect Architecture* with an *Xcelerated Memory Technology* to eliminate bottlenecks, and deliver 15% faster memory throughput, using 667 MHz memory, to access data faster for business critical, time-sensitive transactions. The Xcelerated Memory Technology enables data center staff to install up to eight DIMMs per socket on the x3755 and 6 DIMMs per socket on the x345. Without this value-add technology, competitive systems are limited to four DIMMs, or may be limited to just 533MHz when more than four DIMMs per socket are installed. IBM's *eXtended I/O*, providing 33% greater expandability, slotless RAID, and support for HTX, PCI Express, and PCI-X.

IBM's new blade technology is truly unique, adapting AMD's *HyperTransport* technology with "snap-in" scalability to upgrade an LS41 blade server from two sockets to four (4 to 8 processors, 8 to 16 threads), with unparalleled ease. The *LS41* enables the CIO to protect the enterprise investment in BladeCenter with a "pay-as-you-grow" implementation, eliminating any requirement for over-provisioning of unneeded resources and lowering deployment cost. The LS21 and LS41 also protect the investment made in previous BladeCenter infrastructure.

Four-socket scalability is also available from the rack-mounted x3755. The x3755 is

[2] Double Data Rate 2.

configurable with a pass-through module in lieu of a fourth processor module. The pass-through module will improve inter-processor communication for memory intensive applications and may actually deliver better performance than competitive systems can offer with four sockets.

Perhaps the most significant differentiator for System x to other x86 platforms is the availability of *CoolBlue*, a portfolio of IBM innovations made up of technologies to dramatically improve power utilization and reduce energy costs. CoolBlue includes comprehensive hardware and systems management tools that began to appear in IBM systems in 2004. They allow the data center to optimize power consumption, management, and the cooling of data center infrastructure. See Exhibit 2, below.

Not insignificantly, IBM also announced the availability of the *BladeCenter Migration Factory* to provide human resources and skills that may no longer exist in the enterprise data center. The Migration Factory provides migration planning, tools, and services, allowing the data center to reduce TCO by employing IBM intellectual capital.

Exhibit 2 – CoolBlue Components

- **IBM PowerExecutive** – Allows the data center staff to monitor actual power usage and heat emissions, and cap the power used by any given server;
- **IBM Thermal Diagnostics** – The Industry's first thermal analyzer to identify and take action on heat-related issues in the datacenter;
- **Rear Door Heat eXchanger** – A water-cooled door, using chilled water from the A/C system, to cool the air emitted from the rack, reducing heat emissions by up to 55%;
- **Calibrated Vectored Cooling** – Engineers the path of cool air flow through the system, protecting components inside the server;
- **Power Configurator** – A tool to help the staff plan for data center power needs, delivering better sizing info for System x and BladeCenter.

Source: IBM

Conclusion

IBM continues to introduce open systems servers based upon commodity microprocessors but surrounded by systems innovation to improve data center operations and business efficiency. IBM and AMD have established a viable working relationship where both parties are learning from each other and the enterprise community comes away the winner.

System x, with the latest Opteron technology has placed itself in a role of server leadership with innovations in scalability, I/O, and memory. However, it is in the data center, where energy consumption has had the largest impact on IT budgets, that IBM has made the most significant improvements. With CoolBlue power management, System x can improve the TCO picture of any enterprise.

Moreover, IBM and AMD are clearly not done, having signed an agreement to extend their partnership to 2011. The enterprise data center can look forward to more processor innovation in 45nm, 32nm, and 22nm technology generations to come.

You can acquire commodity x86 technology from many sources, but if you are looking for innovative ways to lower your costs and increase your performance, IBM's System x would appear to have the green light.

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

David Reine **is Director, Enterprise Systems for The Clipper Group.** Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. He joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- ***Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 123. (Please dial "123" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, **The Clipper Group Explorer**, **The Clipper Group Observer**, **The Clipper Group *Captain's Log***, **The Clipper Group Voyager**, and *"clipper.com"* are trademarks of The Clipper Group, Inc., and the clipper ship drawings, *"Navigating Information Technology Horizons"*, and *"teraproductivity"* are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.