THE CLIPPER GROUP
Navigator™

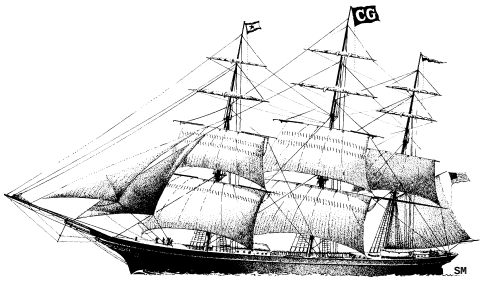

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2006008 February 8, 2006

IBM *BladeCenter H* Improves Performance and Reduces Cost

Analyst: David Reine

Management Summary

It is the middle of the afternoon and you're sitting in your office with a gnawing pain that can only mean one thing: lunch is far in your rear-view mirror and dinner is still hours away. You are hungry and need a snack. If you were at the mall you could stroll over to *Krispy Kreme* for a donut, or *Mrs. Fields* for a chocolate chip cookie, or even *Godiva's* for some chocolate candy. However, instead you go to the snack machine down the hall and view your choices, an interesting selection of heterogeneous foods: candy bars from *Hershey's*, snack pies from *Hostess*, chewing gum from *Wrigley's*, chips, and more. No cashier required, a single slot takes your money and gives you change, depending upon what dynamic choice you make. Potato chips or *Smartfood*, high-fat content or a healthy snack; it is your choice. Once a week, a serviceman will come by, remove the cash and refill the slots that were emptied, replacing what didn't sell with a more popular choice. *One machine with dozens of choices, all controlled by a single control mechanism.*

The information technology (IT) industry has adopted this concept and taken one giant step forward with bladed architectures. Utilizing a concept of shared resources, multiple blade drawers can be installed with heterogeneous processors supported by a common set of I/O controllers, storage devices, and switches, in a single rack, where electrical power and administration can be consolidated to lower the total cost of ownership (TCO) of IT resources for the enterprise. Through virtualization, a single server can be divided into multiple partitions to better utilize a resource that would otherwise be wasted. Blade architectures reduce the TCO of the IT infrastructure, not only through consolidation and virtualization, but also through the economies gained by reducing the electrical requirements to drive enterprise IT and cool the data center. Blade architectures also save in terms of floor space, enabling the enterprise to expand their processing capability without increasing the size of the data center.

IBM has advanced its blade technology even further. With IBM *eServer BladeCenter H*, IBM has introduced an integrated, heterogeneous platform with higher performance and continued power and cooling advantages. With a 10Gb I/O infrastructure, *BladeCenter* can now execute *AIX*, *Linux*, *Solaris*, and *Windows* applications faster, providing mission-critical business solutions with the flexibility to reduce administrative costs by a single point of management. With new blades, the *JS21* and *HS20*, the data center can improve performance and reduce wasted resources. With IBM innovation in cooling through *cool blue*, the staff can reduce air conditioning costs as well. Pretty cool! To learn more about IBM BladeCenter, please read on.

IN THIS ISSUE

The Data Center Dilemma

As we enter 2006, we feel the effects of the energy crisis in our daily life every time we have to heat our homes or fill our gas tanks. As every CFO or CIO knows, we also feel the effects at work. The proliferation of servers has made the cost of running a data center prohibitive from both the personnel and electronic resource perspective. The data center staff has been bloated with an influx of administrators to manage the myriad servers running applications dedicated to specific jobs and specific server environments. There are *Windows* servers here, to manage the email. There are *Linux* servers there, to manage the web services. There are *UNIX* servers everywhere to run mission-critical applications. Moreover, it is not just one of each; there are dozens in every SMB and hundreds in every *Fortune 500* enterprise, each drawing on ever more expensive kilowatts of energy to both power and cool the environment. Every CIO is looking for a scalable business solution to better utilize both the human and electronic resources within the data center and thus lower the TCO of the IT infrastructure. Unfortunately, s/he must be able to improve performance at the same time that s/he is lowering costs.

The data center staff cannot accomplish this by merely trying to consolidate a couple of Window's applications onto one server. The server sprawl that has filled every available square foot of space with mono- and dual-processor servers has necessitated an extreme makeover. The data center needs to be simplified to reduce administrative overhead, as well as consolidated to reduce server count. Servers need to have the flexibility to not only run multiple homogeneous applications within a single CPU, but also the ability the run heterogeneous applications within the same server. The servers need to have virtualization services to better utilize their environments, enabling a single processor to run UNIX, Windows, and Linux applications within different partitions of the same CPU. These servers need to take advantage of the newest multi-core processor innovations to reduce the amount of energy required to power and cool them.

Consolidation of storage resources over the past decade has reduced the number of direct-attached connections to the data center data arrays, the lifeblood of the enterprise. Unfortunately, it has not reduced the number of connections to the resulting storage area network (SAN). A proliferation of servers fills the data center, all trying to access the same information and communicate that information to other servers within the enterprise or to partners. Higher-performing I/O is required. In order to be able to take advantage of the current generation of IT, data centers have been upgrading their architecture from standalone servers, each with their own power source and administration requirements, to a blade environment, where multiple processing boards take advantage of a common infrastructure.

The Blade Environment

Blade server environments were designed to provide a shift in the paradigm away from traditional 1U and 2U rack-mounted servers. Originally designed to lower cost and ease management of the server network, blades were optimized to maximize rack density at the expense of performance. The architecture has evolved to include the latest x64 processors from Intel (*Xeon*) and AMD (*Opteron*) for Windows and Linux, along with commodity RISC processors for Linux and UNIX from companies such as IBM (*POWERPC*) and Intel (*Itanium*), with mono-, dual-, and quad-processor configurations now available. Blade servers now focus on improving virtualization and deployment.

Early blade implementations emphasized consolidation of existing environments. They are now being used in a multitude of capacities, in the data center and remotely, to deploy new applications, no matter what the operating environment. Blades can be dynamically allocated to provide processing power to the business problem at hand, with no concern about having to over-provision rack-mounted servers, in order to handle anticipated peak processing loads.

Due to their compact design, blades did not have the configurability or high-availability features of their rack-mounted predecessors, but did maintain compatibility with industry standards. This commitment to

standards has seen the blade community evolve and organize their efforts into the formation of Blade.org, a collaborative effort of developers in order to accelerate blade solutions and create greater value for blade installations. This effort has lead to the availability of higher performance blade servers, with high-speed networking and the capability to support the demands of the most extreme mission-critical applications for data intensive environments.

IBM BladeCenter

IBM was in the forefront of the shift to blade computing. They introduced *BladeCenter* in November, 2002. BladeCenter is a 7U chassis, supporting 14 blades, mounted vertically to provide for the highest density possible, with up to 6 chassis in a rack. It had mainstream applications for both the enterprise and SMB markets, with configurations ideal for both the data center and remote locations. This was followed by the announcement of the *BladeCenter T* in April, 2004, an 8U chassis configured with 8 blades. The chassis was ruggedized and found applications in the telecommunications and military arenas. Now, IBM is adding to those implementations with a new, high-performance blade solution, *BladeCenter H.*

BladeCenter H has a 9U, 14-blade configuration, designed to support a 10Gb I/O interface. There are up to four 10Gb channels to each blade, with initial availability at a 4X InfiniBand level, jointly developed by IBM and Cisco, a Blade.org partner. There are eight channels available to each blade. This will extend support of blade architecture into data warehousing and commercial analytics, as well as critical applications in the life sciences arena, environments that demand high-performance computing and extreme I/O capability. This enables the data center to integrate blade infrastructure deeper into the IT environment, supporting the most important data intensive applications.

BladeCenter H supports the same blade format as the existing family of products and so every blade and switch can be moved into BladeCenter H. IBM provides support for a wide range of blade options to support any menu of commodity operating environments and satisfy the performance appetite of any mission-critical application. IBM offers an *HS20* Intel Xeon blade and an *LS20* AMD Opteron blade, both running Windows and Linux, and the *JS20 PowerPC 970* blade supporting both *AIX 5L* and Linux. IBM has

Exhibit 1 – Read View of BladeCenter H

A. Advanced Management Module
B. Switch Module Bays 1 & 2 (dedicated ethernet)
C. Switch Module Bays 3 & 4 –OR- Bridge Slots
D. Dedicated Bridge Module Slots (up to 4)
E. Blowers (2)
F. High Speed Switch Module Bays (4)
G. Light Path Diagnostic Panel
H. Aggregated Serial Connector
I. Consolidated Redundant Power Inputs (2)

Source: IBM

now announced the availability of three new blades: a new HS20 with ultra low voltage requirements, a new *JS21 PowerPC 970MP* blade, the first blade with built-in virtualization, and a new *Cell Broadband Engine*-based blade for graphic intensive, numeric applications.

JS21 Blade

The JS21 is optimized for high speed with up to three times the performance of the previous generation JS20. It is available in a two-socket single-core model running at 2.7GHz, or a dual-core model running at 2.5GHz. The JS21 supports up to 16GB of memory, twice the capacity of the JS20, and up to two small form factor (SFF) SAS drives configured for RAID 0 or 1. Supporting up to 10 virtual partitions per core with the inclusion of *Advanced POWER Virtualization*, the same virtualization technology supported on the more high-end IBM system *p5* servers, the JS21 is one of the densest yet cost-effective virtualization solutions on the market targeted at AIX and Linux servers and workload migration. Having achieved the #1 Linpack 2-core and 4-core results among all RISC servers, blade or rack, the JS21 is no slouch for performance either, delivering the goods for high-performance computing clusters specializing in bioinformatics, seismic analysis, and fluid dynamics. The JS21 supports AIX5L, the Red Hat and SuSE versions of Linux, as well as IBM's *Cluster System Manager*. A variety of application solutions will be available to ship with the JS21 blade, including virtualization for server consolidation and a Grid starter kit. The two socket, single-core configuration is priced starting at $2,499, while the dual-core implementation is priced at $3,999. This compares very favorably to other UNIX solutions on a price/performance basis; however, performance will vary depending upon the application. The JS21 is supported in all three BladeCenter chassis.

HS20 Blade

The new HS20 blade is a dual-core, dual-socket low-voltage 32-bit Xeon implementation at 1.67 and 2.0GHz. It is optimized around its outstanding power and cooling attributes and should fit in very nicely in environments still running 32-bit applications. The low voltage Xeon CPU consumes only 31 Watts of power per socket which compares nicely with the AMD Opteron at 68 Watts, for a 64-bit chip, and has outstanding performance/watt characteristics to help lower the TCO. The HS20 supports up to 16GB of memory with two SFF drives, either 36 or 73GB, with RAID 0 and 1 support. The new HS20 is supported in all three BladeCenter chassis.

Cell BE Blade

The third new blade being introduced by IBM is a BladeCenter Cell-based blade. When discussing the newest high-performance, multi-core technology, the conversation must begin with the *Cell Blade with Broadband Engine (BE)*. Designed with a breakthrough architecture employing a nine-core processor, the BE blade is targeted at high performance workloads for digital media, medical imaging, aerospace and defense, and the communications industries. With a scheduled availability of 3Q06, the Cell BE employs one *Power Processing Element* (*PPE*) and eight *Synergistic Processing Units* (*SPU*) connected via a high speed data ring implemented for a throughput of 192 GB/second. This is referred to as the *Element Interconnect Bus (EIB)*. The EIB extends transparently across a high-speed coherent interface between the dual Cell BE processors, running at 20GB/sec in each direction. The BE is a double-wide blade, with up to seven blades per chassis, with one IDE drive supported on each blade. The BE blade also has two embedded 1Gb NIC interfaces and two InfiniBand daughter cards on each for connection to external I/O. IBM has made evaluation software for the Cell blade available on their *Alphaworks* website and Open Source software at the University of Barcelona website.

BladeCenter H also has a customer-removable and hot-swappable media tray that connects directly to the chassis mid-plane. The BladeCenter H chassis is powered by four 2900W power supplies and is designed for ease of maintenance with easily-accessible components and a full *Light Path Diagnostic* panel to facilitate maintenance. (See Exhibit 1, on the previous page.) An innovation from IBM is its unique thermal design that

Exhibit 2 – Advanced Management Module Features in Initial Release

- Full compatibility with Management Module;
- Granular role-based management permissions;
- Automated network failover for the module;
- Automated network uplink failover;
- *PowerExecutive* reports actual power consumed by each blade;
- SMASH CLP proxy for Linux/Unix users;
- SNMP configuration changes without restart;
- First Failure Data Capture.

consumes 20% of the power of traditional fan systems and uses shared components to reduce the number of power-consuming parts, all contributing to a reduction in TCO. In addition the blade rack can be configured, optionally, with IBM's *Cool Blue* cooling environment to further contribute to significant energy savings.

Using technology derived from the mainframe service processor, IBM has implemented an *Advanced Management Module* for BladeCenter H, and it is also backward compatible with the standard BladeCenter[1]. This module manages BladeCenter H at the chassis level, not the server level. A simplified management system allows a single administrator to manage more through automation, with additional features planned for the future, including rack-level management to further simplify blade control. The BladeCenter H advanced management module supports all of the features of its predecessor, plus the features described in Exhibit 2, above.

Conclusion

Information Technology has seen a significant paradigm shift from what we used to perceive as important. Only a few years ago raw processing power was the goal and the power to run it was just a commodity. Now the processing capacity is the commodity and reducing the power consumption is the goal. These recurring costs are the key to reducing total cost of ownership. Blade architectures enable these savings while maintaining or improving performance levels.

BladeCenter can save over 20% of your energy costs through savings from power and cooling, and in the space you save in your data center, and it does that while enabling a heterogeneous environment in which the data center has access to applications written for any of the leading commodity operating environments. Furthermore, with BladeCenter you also have access to the latest processor technology from Intel and AMD, as well as IBM. If you are considering a blade environment for your enterprise, take a look at IBM's offering; it may be the best option for you.

[1] BladeCenter T uses a different form factor but runs the same code base.

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.