



IBM Increases Blade Performance, Lowers Cost — Eliminating the Fight, Enabling the Switch(es)

Analyst: David Reine

Management Summary

Traffic is one of the daily headaches that we have all been forced to live with. For some, traffic means driving along slow, two-lane roads, negotiating a never-ending series of red lights every other block. For others, it means navigating a high-speed, four-lane turnpike, speeding the commuter to a row of tollbooths, and long delays, awaiting every few miles. Some states have implemented a fee-based wireless recognition system, called *FastLane* or *EasyPass*, which enables the toll collection system to identify a transponder on the auto windshield and collect the toll electronically in an express lane. This simplifies the process and enables the commuter to bypass long lines, continuing through the plaza without stopping. A charge for the transponder creates an elite group of commuters willing to pay an additional fee to save time, and perhaps even more money, in order to gain quicker access to a high-speed thoroughfare to their destination. Other states have distributed transponders at no charge in order to simplify further the collection process, eliminating more toll collectors and lowering the overall infrastructure cost.

Similarly, to a 21st century enterprise, traffic refers to a complex communications process and the need to gain access to critical, high-speed I/O paths. Years of uncontrolled expansion have led to a nightmare of server proliferation within the data center. Underutilized servers, each running a single application, result in an enormous waste of resources, both human and natural. The cost to manage thousands of mono- and dual-processor servers eats away at the IT budget. The inconvenient truth for the enterprise is the need to lower rising temperatures in the data center caused by the energy wasted in running the IT infrastructure. Reversing the warming of the data center can enhance enterprise profits. In order to reduce the total cost of ownership (TCO) of the IT infrastructure, the CIO must consolidate these disparate servers into a simplified, manageable system, with many applications sharing compute resources in a virtualized environment, and a workload management system making the most efficient utilization of the IT environment. Blade architecture has become the hottest method to cool the data center, with IBM leading the way in developing a flexible, yet simplified, architecture to share resources and move applications to underutilized processors. As faster multi-core processors appear, the need for improved communications between processors, as well as to the network, becomes more critical.

With *BladeCenter H*, IBM has enabled the enterprise with the capability to consolidate a wide-range of applications and share a common set of resources, including the latest in high-speed I/O technology. To see how *BladeCenter H* can help your data center, please read on.

IN THIS ISSUE

➤ The Data Center Environment	2
➤ Blade on the Cutting Edge	2
➤ IBM BladeCenter H	3
➤ Conclusion	4

The Data Center Environment

Today's enterprise data center faces a long list of pain points as it transitions from an era of cost cutting to stay afloat while meeting minimum IT requirements, to a period of expansion and growth. How can the next generation data center maximize performance *and* reduce the cost of running a mission-critical application set in an enterprise populated with a heterogeneous mix of under-utilized servers running open system applications under *Windows* and *Linux* and mission-critical applications under UNIX? How can they do this, while, at the same time, the cost of energy rises at a precipitous rate? The average x86 server is now estimated to operate at a processor utilization rate of only about 15-20% and the data center staff has no good way to dynamically (or automatically) deploy new applications or services, or re-provision their tower and rack-mounted servers to take advantage of the unused resources. Unfortunately, these underutilized platforms do not consume *only* 15% of the power required to run them and to cool the data center from the heat they generate. In most cases, their power supplies consume 100% of their rated energy requirement – **wasting up to 85% of their electricity and a significant portion of your IT budget!** In response, many IT organizations have begun to implement consolidation programs, merging multiple, older x86 servers onto fewer new platforms utilizing the new multi-core architectures (often at up to 10:1, or greater) in order to improve business processes. Virtualizing the environment (scale-in) enables the enterprise to run multiple applications over shared processing resources. With consolidation and sharing the workload, the processing power of a single CPU can cure many IT pains by:

- Significantly improving server utilization;
- Reducing the number of servers required;
- Reducing the kW used to power the data center;
- Reducing the energy required to cool the data center, lowering utility overhead; and
- Reducing management costs to administer excess, over-provisioned platforms, improving productivity.

Unfortunately, virtualization, multi-core technology, and consolidation are all driving an increased communications workload and increasing additional port requirements.

Once the decision has been reached to redesign the enterprise IT environment, the CIO and data center staff must agree upon an appropriate computing environment. Do you want to install a *scale-up* platform, typically running mission-critical UNIX applications with very high memory

scalability? Do you go with a *scale-out* architecture, typically deployed on 1-way and 2-way, Windows and Linux platforms running infrastructure solutions? Or do you try a third approach, implementing a *scale-in* environment, using virtualization tools to share all resources with a blade architecture.

Blade on the Cutting Edge

Blades represent a multi-billion dollar market, with a growth rate of almost 10% a year, a clear indication that blades are making significant inroads on the rack-mount business.¹ Acquisition pricing continues to fall on a price/performance basis as a new generation of multi-core processors from AMD, Intel, IBM, and Sun provide better performance and energy savings. Blade vendors need look no further than commercial razor blade companies for a successful marketing strategy. Companies give away their razors (i.e., sell them cheaply) in order to sell *their* blades. Discounts on racks and chassis and buyback programs could be imminent as competition heats up and vendors attempt to gain a blade foothold in the data center.

Blades provide the simplest path for the IT staff to do more with less. Recent advancements in processor technology, such as multi-core CPUs and virtualization, have on the one hand, enabled data centers to lower application-licensing costs, while also creating a need for an even more performant I/O architecture. The new I/O environment must have a greater bandwidth to enable the transition of applications from one processing environment to another, with the capability to support the higher bandwidths necessary for new consumer technologies such as IPTV and on-line gaming. The new enterprise IT environment needs to include support for 10Gb Ethernet (e-Net) and 4Gb Fibre Channel (FC) SAN I/O protocols, among others, and a high-speed switch technology to enable the on-demand scalability required in blade architectures. Streamlined blade designs enable the enterprise to reduce both capital expenditures and operating expenses, with a blade environment typically using 30% less power than comparable rack-mounted servers due to centralized I/O and power resources. Blades also provide flexibility, allowing the sharing of bandwidth between blades. This enables a denser installation and a cooler data center environment. Blades simplify the implementation, deployment, main-

¹ Remember that each of the new processors may be ten times as powerful as the ones that are being consolidated. If the "horsepower" of the new servers were the metric of growth, the rate would be many times larger.

tenance, and modular growth of the IT network, at the same time improving administrative productivity. Blades also enable the enterprise to manage server-provisioning better, while improving business continuity at a price significantly lower than that of a fault-tolerant platform.

Blade architecture allows the enterprise to consolidate the entire data center environment, not just server processing. This includes sharing I/O and storage resources – enabling the enterprise to solve real business problems. IBM initially announced a blade architecture in 2002 with *BladeCenter*. Recent introductions from IBM in microprocessor technology, energy efficiency, and server-surround architecture have led to the introduction of a second generation of blade servers, the *BladeCenter H*.

IBM BladeCenter H

In 2006, IBM refreshed their BladeCenter architecture by introducing BladeCenter H with ten times the I/O performance, twice the I/O connectivity, and twice the I/O capacity. They also introduced BladeCenter H with the latest x86-microprocessor technology from both AMD and Intel, with the dual-processor *HS21*, based upon dual- and quad-core *Xeon* processors from Intel, and the dual-processor *LS21* and quad-processor *LS41*, based upon dual-core *Opteron* processors from AMD. These blades enable the enterprise to take advantage of the high-performance, low energy requirements of these commodity processors to consolidate *Windows*, *Linux*, *AIX*, and *Solaris 10* infrastructure applications and lower the TCO of the data center. In addition, IBM has also refreshed their *POWER* blade product set, with the *JS21*, a dual-processor blade based upon the dual- or quad-core *POWERPC 970MP* architecture. This blade enables IBM to extend the scope of BladeCenter H to mission-critical applications written for *AIX* and *Linux* on IBM's *System p* servers. Adding IBM's *QS20* blade with two *Cell BE* CPUs to the mix, gives you a blade environment with the widest choice of processors in the industry

In addition to multi-core processors that are a key to consolidation, BladeCenter H takes advantage of a wide range of system features to lower its energy requirements, from its basic design to its virtualization and advanced cooling technology.

- IBM designed BladeCenter H for better energy efficiency. By sharing system infrastructure components, and using more efficient power supplies, IBM reduces power consumption, by up to 30%, as compared to other architectures.

- *Calibrated Vector Cooling* is another energy-saving technology, using dual air-paths to improve reliability. Combined with tools such as *Power Calculator* and *Power Executive*, BladeCenter H can reduce the cost of the power infrastructure.
- By virtualizing the processor environment into a blade architecture, the data center can implement multiple applications on each processor core, retiring a great number of the mono-core, rack-mounted servers that have been driving the enterprise. This will save energy in running the applications, and in cooling the environment. IBM also has implemented their *CoolBlue*² technology for BladeCenter H further reducing the amount of heat released into the data center.

IBM customers have seen a continuous stream of improvements to the BladeCenter architecture, not only in server blade design, but also in I/O processing performance. Originally announced with 1Gb E-Net support, the BladeCenter has seen an evolution in I/O capabilities, with 2Gb F.C. in 2003, to 1X InfiniBand support in 2004, and 4Gb F.C. support in 2005. Now, IBM has extended the I/O capabilities of BladeCenter H through partnership and collaboration with some of the leading communications companies in the IT industry. IBM has created *Virtual Fabric Architecture (VFA)* with support for hundreds of servers. With VFA, BladeCenter H has double the number of ports per blade and a choice of the fastest technologies available to increase the speed and efficiency of data transfer in applications such as IPTV and online gaming.

10Gbps Ethernet Switch

Working in partnership with **Blade Network Technologies** (BLADE)³ and **NetXen**, IBM has developed a new 10Gb E-Net solution consisting of a 10Gb switch from BLADE and a dual 10Gb port I/O adapter from NetXen, specifically for BladeCenter H, to satisfy a need for higher bandwidth. The switch module has twenty 10Gb ports, with 14 interfacing to blades enabled with an expansion card, and 6 XFP-based connections to the network. This manifests as a significant increase in 10Gb throughput efficiency and a comparable reduction in 10Gb latency.

This is the first solution with full 10Gb capability. With a list price of \$9,799, the switch costs

² See **The Clipper Group Navigator** dated August 22, 2005, entitled *IBM Makes Big Blue "Green" – Using Open Systems to reduce Data Center TCO*, available at <http://www.clipper.com/research/TCG2006073.pdf>

³ A spin-off from Nortel.

less than \$500 per port. This represents a significant savings when deploying a bladed 10Gb E-Net network versus a rack-mounted solution⁴. This switch will be available in February.

IBM InfiniBand Switch Options

InfiniBand provides a unique capability to both unify and simplify the communications infrastructure in the data center, while lowering cost and increasing performance. BladeCenter H supports two InfiniBand offerings, a switch developed in collaboration with **Cisco**, and a pass-thru module in partnership with **Voltaire**, both available today. The Cisco solution is a 4X (10Gb) switch providing dual 4X connectivity to the blade, with 80Gb of bandwidth to the network. Cisco's *VFrame* software virtualizes the I/O for a scalable, unified server environment. The Voltaire solution has 14 non-blocking 4X DDR connections to each blade, providing up to 280Gb of pass-through interconnect through the chassis.

Cisco Fibre Channel Switch

IBM has introduced a new 4Gb Fibre Channel (FC) switch from Cisco, giving the data center staff another option. BladeCenter H already supports FC switches from QLogic, McData, and Brocade. Now, the data center can install an end-to-end Cisco network.

InfiniBand Bridge Modules

In conjunction with QLogic⁵, IBM is introducing two new InfiniBand bridge modules, with January 2007 availability, to help reduce data center complexity by further integrating functionality into the chassis: one for E-Net, and one for FC

- The Ethernet module provides six RJ-45 connections to an external data network, with two internal 4X ports to the InfiniBand Switch. The external ports support 10/100/1000Mbps.
- The F.C. Bridge Module supports six SFP-based ports to an external SAN and two internal 4X ports to the InfiniBand Switch. The bridge supports one, two, and four Gbps SANs.

Multi-Switch Interconnect Module

IBM designed BladeCenter H to support eight I/O paths per blade. Furthermore, each of the new blades certified in the system, HS21, LS21, and JS21, have eight paths with the LS41 providing twelve paths. Configurable in a high-speed switch slot, the *Multi-Switch Interconnect Module (MSIM)* enables each blade to exploit that capability for each of the switches in the BladeCenter H port-

folio. The MSIM also continues the flexibility of the BladeCenter H with a variety of interconnection configurations, including:

- Six 1Gb E-Net and two 4G F.C.
- Four 1Gb E-Net and four F.C. and
- Two 1Gb E-Net, two 10G E-Net, and two 4G FC.

BladeCenter Address Manager

To simplify the deployment and management of blades, and as a part of VFA, IBM plans to deliver in 2H07 *BladeCenter Address Manager* to virtualize E-net and FC connections within a system. This utility assigns port addresses to the blades via software, allowing the data center to give specific LAN and SAN assignments to the blades for fast repurposing of the servers. This also decreases the time required to provision new blades. Additionally, BladeCenter H is easy to manage with self-discovery and an intuitive GUI.

Conclusion

With BladeCenter H, IBM addresses all major pain points facing the enterprise CIO.

- They have solved the problem of server sprawl with the capability of installing high-performance, multi-core blades;
- They have taken the complexity out of a heterogeneous environment with the widest array of blade servers in the industry;
- They have attacked the total cost of ownership problem with an innovative program of energy-savings technology, from low-energy blades to CoolBlue; and
- They have simplified the interconnect maze with a complete set of I/O components from the leading E-net, FC, and InfiniBand vendors.

IBM continues to lead the way in blade architecture by delivering innovative technology while remaining compatible with the original BladeCenter design, protecting the enterprise's investment in previous BladeCenter technology.

BladeCenter H provides the data center with the widest offering of open systems architectures – choice of processor, choice in storage, choice in I/O - and enables the survival of the coolest. BladeCenter allows the data center to choose the configuration best suited to its enterprise environment. If your enterprise is looking for choice, look to BladeCenter H. You may find your answer.



⁴ The dollar value of the savings is configuration dependent.

⁵ Through the acquisition of SilverStorm.

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, Clipper Notes, 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.