THE CLIPPER GROUP
Navigator™

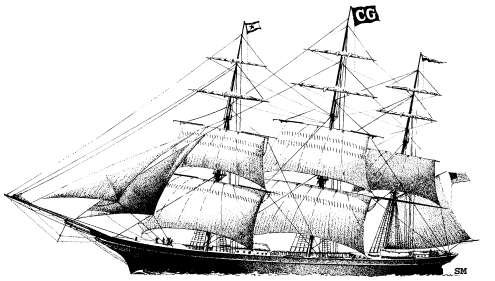

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2008053 October 10, 2008

Improving Performance, Scalability with Standards — IBM One-Ups the Competition

Analyst: David Reine

Management Summary

Standards enhance and simplify our lives at a very basic level. In the world of entertainment, the Motion Picture Association of America (MPAA) has taken a complex issue, what is appropriate for our children, and simplified it down to a set of five ratings, from *G* for viewing by anyone, to *NC-17*, restricting viewing to those over 17. We know from the ratings that are assigned by the Rating Board whether the movie has elements of sex, nudity, drugs, or violence to which we do not want our children to be exposed. Standards also assist us in shopping for food products, where the government establishes grading systems for the quality and sizing of a variety of items, such as meat and eggs, enabling us to distinguish between *Prime*, *Choice*, or *Select* (among others) for meat, and *Large*, *Extra Large*, and even *Jumbo* for eggs.

Standards have had a significant impact within the data center of every enterprise, also. While many enterprises continue to use mainframes and high-end *UNIX* servers to run the mission- and business-critical applications that drive corporate profitability, a great many more have migrated to an open systems platform based upon Intel's *x86 EM64T* architecture. Because of the widespread acceptance and availability of this environment, the low cost of servers based upon Intel's *Xeon* or AMD's *Opteron* microprocessor have made application platform purchases simpler *and, yet, more complex* at the same time. Because the IT staff knows that any application written for the x86-instruction set will run on any Xeon or Opteron CPU, they know that they have a wide variety of servers from which to choose. Further, with a high degree of competition, the pricing for these commodity platforms, on a price/performance basis, is continually dropping. However, even though each of these open systems platforms is based upon the same, or similar, microprocessors, the performance that they can achieve varies widely, based upon the architecture that the manufacturer chooses to integrate around that standard CPU. For, while standards ensure compatibility, manufacturer's differentiators will determine important factors such as performance, handling, and durability. The choices are endless, ranging from the number of sockets on the motherboard, to the number of cores within each microprocessor, to the chipset that connects the compute engine to the rest of the server, to the memory subset. AMD even chose to incorporate the memory management component within the Opteron CPU itself. *How does an educated IT staff determine which open systems architecture is best?*

IBM has simplified that decision for the data center, with the introduction of a second version of the IBM *System x3850 M2* and the *x3950 M2*. IBM is delivering to the enterprise an open, yet superior, platform for the execution of database services, consolidation through virtualization, and ERP applications. To learn more about the newest System x servers, please read on.

IN THIS ISSUE

Enterprise Application Requirements

A virtual plethora of *under-utilized*, scale-out servers has taken over the data centers of almost every enterprise. If a hacker plants a single virus on just one mission-critical enterprise server, disabling 80% of its compute cycles, the IT staff will employ a strike team to seek and destroy that virus and restore the server to full operating status. How do we rationalize, then, the conditions that exist today where many servers, configured in a scale-out environment, are operating at only 15 to 20% efficiency, while the IT staff ignores the rest of the CPU's capability? The cost to power these under-utilized servers, and cool the data center, increases the TCO of the data center and has a direct impact on profitability. In addition, wasting natural resources limits the amount of energy available to run new applications required to meet the needs of a growing enterprise. *Consolidating* multiple servers onto a single scale-up[1] server, while *virtualizing* the environment to enable multiple applications to operate in parallel is one method of removing much of the complexity from the IT architecture and simplifying the infrastructure, enabling the staff to restore order to the data center. **While it is painful to think about redesigning the IT environment, it is irresponsible to ignore the necessity**.

Improving application performance is always one of the primary objectives of the CIO. However, two of the most important considerations facing the CIO in changing the compute paradigm of the data center are:

- **Meeting the need to make the architecture more energy-efficient**, reducing enterprise demand on electricity, thereby improving performance per watt and making the data center environment more green; and
- **Reducing the total cost of ownership of the data center architecture**, while improving the reliability, availability, and serviceability (RAS) of the application platform.

The data center acquired many of the servers installed throughout the enterprise in the past 3-5 years, in the value-packed era immediately following the tech-bubble bursting in the early part of this decade. These servers were single-socket primarily, single-core processor platforms from companies using the x86 architecture, with microprocessors from Intel and AMD, in scale-out environments running *Windows* and *Linux* applications. They were not designed for the scalability available in SMP architectures used in many *UNIX* platforms, designed to grow with the enterprise. While lacking many of the RAS characteristics found in UNIX and mainframes systems, these x86 scale-out platforms were characterized as "good-enough" for the infrastructure environments and applications that they ran. This may be true for a print or file server, but when a data center installs hundreds of platforms in a mission-critical environment, good enough usually is *not* good enough. **Consolidating these disparate systems requires the highest levels of reliable, scalable performance and I/O throughput, not just *good enough*.**

Over the past few years, there has been an increased emphasis to deliver systems to satisfy the requirements and desires of the most common denominator, the small and medium business, the SMB. With a typical user count of less than 1000, there are literally thousands of SMBs with a requirement for anywhere from 10 to 50 servers to deploy in a scale-out, rack-mounted architecture, or packaged in blade chassis, to satisfy their mission- and business-critical application needs. Now, there is a striking need to apply that same innovation to the enterprise data center where legacy applications are outgrowing the scalability of their current mission-critical platforms. In order to accomplish this, enterprise data centers are now looking to deploy high-end, multi-socket, SMP platforms implemented with multi-core processors to achieve the desired scalability and virtualization to improve server utilization and reduce the TCO.

In fact, multi-core processors have now taken over the mission-critical server domain, not just Intel's *Xeon* or AMD's *Opteron*, but from legacy vendors who have been building scale-up servers for two decades: IBM with *POWER*, HP with *HP-UX*, now on *Itanium*, and Sun and Fujitsu with their *SPARC* architecture. Data centers in this domain often are running workloads such as database servers, server consolidation using virtualization services, and ERP applications. Both back office and high-performance computing applications have a requirement for high scalability and a large shared memory. They also share a need for high availability, improved response times, and improved utilization, as well as a requirement to protect business- and mission-critical data.

[1] See the issue of Clipper Notes dated September 23, 2008, entitled *Perceiving the Dark Side of the Moon – Knowing When Scale-up Computing Makes Sense*, and available at http://www.clipper.com/research/TCG2008048.pdf.

In all cases, these data centers must reduce the IT infrastructure in order to reduce the TCO and improve profitability. They can accomplish this through consolidation and virtualization, reducing the number of platforms, improving the utilization of each platform, lowering energy consumption; and reducing floor space. They cannot, however, lose any of their RAS characteristics. These must remain, or improve, in order to satisfy enterprise service level agreements (SLAs)

By upgrading the existing architecture with denser, more highly scalable platforms, the data center can reduce the TCO without radically changing the compute paradigm. By replacing the existing environment with servers that are binary compatible with prior platforms, the IT staff can reduce application porting costs and staff retraining. Ease of deployment and ease of use is essential. Intel has facilitated this shift with the introduction of the *Xeon MP 7400*, which is faster, denser, and more powerful than its predecessors, and known internally as *Dunnington*, a six-core processor running at up to 2.66GHz with up to 16MB of L3 Cache, at 130W[2], down to 21.67W per core. It is capable of being configured and deployed in multi-socket configurations. The big question that remains is: Whose implementation of the Xeon 7400 is the best? One company staking a claim to that title is IBM with their *System x3850 M2* and *x3950 M2* platforms.

The IBM System x Solution

On September 16th, IBM announced the introduction of the second generation of their System x3850 and x3950 platforms, now known as System x3850 M2 and System x3950 M2. With the Intel 7400 microprocessor as the compute engine, IBM can configure four CPUs, with six cores in each, in a single 4U x3850 M2 chassis. Further, with the addition of a *ScaleXpander* option kit, can protect the enterprise investment in that x3850 and upgrade it to an x3950, capable of scaling up with three optional 4U chassis, to enable access to 16 Xeon 7400 CPUs with 96 cores. **This provides the enterprise with the basic platform to consolidate the IT infrastructure of the data center, virtualizing a variety of applications to fully utilize the processing power and memory capacity of this new System x server.** In fact, with a standard 42U rack, the data center can install 40 CPUs with 240 cores to provide enough processing power to consolidate almost any environment.

Exhibit 1 –
x3950 Active Memory Features

- **IBM Chipkill** – corrects multiple single-bit errors to keep a DIMM from failing for up to 16x better error correction than standard ECC memory;
- **Memory ProteXion** – uses redundant bits in a data packet to provide backup in the event of a DIMM failure;
- **Hot add and hot-swap of memory** – High availability feature to add or replace new DIMMs without turning off the server, thereby increasing the amount of RAM available to the operating system;
- **Memory Mirroring** – Roughly equivalent to RAID-1 in disk arrays, in that memory is divided in two ports and one port is mirrored to the other half;
- **Energy Efficient** – Utilizes 37% lower power requirements than fully buffered DIMMs.

Source: IBM

With its built-in scalability, the x3950 provides the data center with a unique capability to grow from two Xeon processors to 16, as application needs change. There is no need to overprovision the x3950; it can grow as your requirements grow, with *XpandOnDemand* scalability, enabling the data center to minimize acquisition and deployment costs, controlling the TCO of the application platforms, while obtaining a large SMP system at an entry-level price.

While other vendors manufacture platforms with Xeon processors using Intel's standard chipset for interfacing to the Xeon 7400, IBM chose to differentiate System x with their own innovative architecture, the *Enterprise X-Architecture* (*eX4*), currently on its fourth generation. IBM designed eX4 for scalability, and maximum processor, memory, and I/O utilization, boosting performance and reducing operating costs. In addition to supporting processor scalability with four 1066MHz front side buses, the eX4 supports up to 256GB of PC5300 DDR2 SDRAM memory per node[3] , using 32 DIMMS, with outstanding reliability features. (See Exhibit 1, above.).

[2] The 2.13GHz version consumes only 50W.

[3] Support for 1TB per fully configured 4-node system.

All 256GB is available to any of the processors installed in the x3950. The eX4 also supports up to 256MB of L4 cache per node. In addition, eX4 supports:

- Up to seven high-speed PCIe x8 adapter slots per node, two are hot-swappable;
- Up to 587GB of internal storage, on four hot-swap SAS drives, per node with RAID 0/1 protection standard, 2.3TB per 4-node system; and
- The *ServeRAID-MR10k* RAID controller with 256MB of battery-backed cache, for RAID 0, 1, 10, 5, 50, 6, or 60 protection;

By adding IBM innovation into an industry standard platform, IBM has achieved 29% better performance per watt and 13% lower power consumption than the competition[4].

The x3950 M2 also contains a full set of built-in hot-swap features, including disks, memory, I/O slots, power supplies, and temperature-controlled fans with *Calibrated Vector Cooling*, to enable ease-of-use deployment and maintenance, along with the *Remote Supervisor Adapter II* that communicates with the *Baseboard Management Controller* to enable the user to manage and control the server, locally or remotely. All of these features make the x3950 M2 an ideal vehicle for hosting consolidation with virtualization. In fact, selected models of the x3850 now come with *VMware ESXi 3.5* preloaded on a USB 2.0 Flash Key, supporting a diskless configuration with a smaller memory footprint and high performance to simplify deployment of a virtualized environment.

IBM has invested a significant amount of time and resources in creating a differentiated server-surround environment for the standard Intel Xeon architecture. They did this in order to improve the performance, reliability, and ease-of-use of System x over competitive platforms. *How well did they do?*

IBM System x Benchmark Results

Based upon the results of several recently published benchmarks, IBM has achieved an excellent return on its investment. Any enterprise operating in an OLTP environment typically is most concerned about the number of transactions per minute that their platform can execute. In recently published results, the IBM System x 3950 became the first 8-way x86-64 server to exceed 1,000,000 TPM, at **1,200,632**, an increase of 40% over a similarly configured Xeon 7300 platform. In fact, the only two other 8-way servers to exceed the 1M TPM mark both use the IBM *POWER* architecture[5] with the AIX operating system. The significant point here is, at $1.99 per transaction, the x3950 has a price/-performance ratio 43% less than the POWER models ahead of it. No other vendor has even published an 8-way Intel 7400 result, indicating that IBM has a significant edge in scalability.

In the 4-way category, the IBM System x3850 leads with a TPM rating of **684,508**, indicating exceptional scalability with the x3950, taking into consideration that the x3850 was tested with *Windows* and *SQL Server* while the x3950 was tested with *Linux* and *DB2*. The x3850 has a sizable edge in performance over the HP *ProLiant DL580G5* rated at 634,825. As additional benchmarks are performed, the IBM System x3950 M2 can be expected to be found at or near the top.

Conclusion

Once again, IBM has proven that there is value in innovation. Using a leadership design and unmatched open systems price/performance scalability, IBM has delivered a superior pay-as-you-grow platform for a variety of enterprise data center requirements. If you are concerned with consolidation and virtualization, the System x3950 has the scalability and flexibility you need to reduce the TCO of your data center. If the enterprise workload in focused on database performance, IBM has proven, once again, leadership in transactional performance. If your data center is committed to increasing its compute capability while at the same time controlling energy consumption, the System x3950 may be the platform you need.

[4] See *Principled Technologies Report* at http://www.principledtechnologies.com/Clients/Reports/IBM/IBMvCon1p0808.pdf.

[5] The two systems are the IBM *System p570* and the Bull *Escala PL1660R*.

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.