THE CLIPPER GROUP
Navigator™

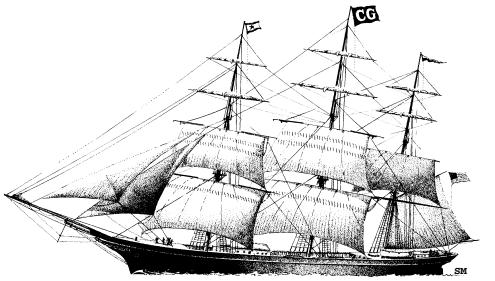

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2009031 July 28, 2009

IBM Refreshes System x Product Set — Introduces 6-Core Opteron to the Data Center

Analyst: David Reine

Management Summary

Americans simply love our cars. We get attached to the features; we get attached to the look and the feel. In fact, we get so attached that we often try to overlook the faults that creep in as the years go by. In an era when the cost of energy is doubling, or even tripling, the gas economy of some of our vehicles simply has become unacceptable. New hybrids are appearing every day, replacing the hulking internal combustion engines that have powered our oversized cars for decades. Not only are we wasting fuel, we are emptying our wallets every time we stop at a gas station. We ignore the fact that modern technology can provide us with some built-in conveniences that were not available three or five years ago, when we purchased our current ride, or simply too expensive. Options such as built-in GPS systems not only can save time and money, but they also eliminate the need to stop and ask for directions (if you can admit to doing that!). Yes, there is an added acquisition cost, however, how do you save in time and money by not getting lost in the first place. More important, is the matter of maintenance. When the new car warranty expires, the cost for any repairs, minor or major, becomes an added economic burden and nuisance that many do not wish to face. As we can see, the total cost of ownership (TCO) of our cars is significantly more than just the acquisition cost.

Total cost of ownership is not simply a concern of the car owner. Every enterprise around the world is concerned about the TCO of the IT infrastructure in its data center as well as in every branch or remote office. Due to the increase in the number and use of applications being deployed, as well as the rapid growth in data being stored, server sprawl has become a serious issue throughout the enterprise. Underutilized servers have multiplied like rabbits throughout the IT environment, wasting valuable processing resources and the natural resources that are being consumed to have power and cool infrastructure in the data center. Many of these servers, like the automobiles, are three-to-five years old, configured with older, now dated technology, typically single-core or dual-core x86 processors, with expired, or expiring, warranties. **Acquisition cost is no longer the most significant component of the TCO of the IT infrastructure; it is merely the initial component**. Energy, maintenance, software licensing, floor space, and administrative staff all factor highly in the TCO equation. Thus, the consolidation and virtualization of data centers has become critical to gaining control of the data center budget.

To provide more and better processing capabilities to better support multiple applications on a single virtualized open server, chip vendors such as AMD and Intel, have been innovating with multi-core technology for several years. The latest development is the entry of a 6-core CPU by AMD, referred to as *Opteron Istanbul*. Leading server vendors are now incorporating the Istanbul processor into their latest server offerings. The newest server platform to appear is the IBM *System x3755*, designed for both flexibility and performance to enable the data center to scale the system as business requirements change. To learn more about how the x3755 can help you achieve your TCO goals, please read on.

IN THIS ISSUE

Server Sprawl in the Data Center

Many of today's enterprise data centers are an aging mix of open systems servers based upon single- and dual-core Intel and AMD processors, based on the industry standard x86 architecture. These servers are often deployed with a single business-critical or infrastructure application, running on seriously under-utilized x86 platforms. As a result, the enterprise I.T. staff must maintain and support, potentially, thousands of servers, procured over many years and spanning processor generations. This server sprawl has created many challenges for the data center staff, adding cost to the data center's operational budget. Thus, simplifying the I.T. environment, improving utilization, and reducing the TCO of the data center, have become major points of focus of the enterprise executive team.

Many enterprises have looked at changing the I.T. environment to simplify the infrastructure, consolidating older open systems onto a network of new, multi-socket, multi-core x86 servers. However, when deployed with a single application, these servers would still only utilize a fraction of the CPU resources available to them; often, this is less than 20% of the server's compute cycles. Not only does this waste I.T. processing resources, but it wastes valuable energy as well, adding significantly to the data center's TCO. Equally important, however, the CIO has to be concerned with retaining the investment that the enterprise has made in their application mix, along with the investment made in support and management software. The problem facing the I.T. staff then becomes: ***How can the data center reduce the TCO without negatively affecting the performance and reliability of the IT infrastructure?***

In order to reverse this trend, simplify the I.T. architecture, and lower data center TCO, the I.T. staff must change the paradigm of the data center by focusing on (1) the consolidation of mission- and business-critical servers, as well as infrastructure servers, and (2) the virtualization of these platforms to improve server utilization. In addition, the data center must be concerned with the implementation of programs to support "green" I.T. – to lower energy consumption and help to reduce the TCO of I.T. infrastructure.

Ideally, the I.T. staff would select a single scalable server architecture that would satisfy the needs of the application set today and allow for expansion, protecting the investment that the enterprise is making. In addition to the data center, the I.T. staff needs to have common server configurations that they can propagate throughout the enterprise, to simplify the deployment and support of the infrastructure.

Normally, the enterprise architecture is a heterogeneous mix of On Line Transaction Processing, database, infrastructure, and High Performance Computing applications, some requiring high levels of I/O, with others demanding the highest computational processing capabilities. In order to improve server utilization, some of these platforms will have to be virtualized, i.e., subdivided to run multiple applications, requiring more memory than many of the CPU-intensive applications.

Within any open architecture, the data center has options. There are servers based upon Intel's *Xeon* architecture and systems based upon AMD's *Opteron* technology. Even within these families, there are options. The *Intel 5500* processor (*Nehalem-EP*) is a quad-core CPU for dual-socket systems, while *Intel's 7400* processor (*Dunnington*) is a 6-core CPU, scalable for four-socket systems. AMD matches these architectures with their quad-core *Opteron* (*Shanghai*) and their newest processor, the 6-core *Opteron* (*Istanbul*), scaling to eight sockets. The total number of cores, which is the number of cores per processor times the number of processors, is one indicator of the power of a server.

What the enterprise with multiple requirements needs is a vendor who can provide the right platform for the right application. With *System x*, IBM can provide scalable systems to meet the data center's needs. In fact, IBM has recently added to their Nehalem-EP System x servers with a new scalable platform based upon AMD's Istanbul technology.

AMD's 6-Core Opteron - Istanbul

AMD has just announced the availability of a 6-core Opteron processor, complementing the quad-core Opteron announced last year. This processor has been developed with 45nm technology for reduced power consumption and an increased frequency to support the data center's need for improved applications, expanded virtualization capabilities, energy efficiency, and better economics resulting in a lower TCO. Designed to use the same power envelope as the 4-core Shanghai processor, Istanbul enables increased capabilities for an HPC environment or increased database performance, while reducing TCO through maintaining the same energy

Exhibit 1 – Opteron Features

- ***HyperTransport Technology (HT) Assist*** – to reduce cache traffic between processors.
- ***Increased HT3 Bandwidth*** – to increase the interconnect rate from 2GT/s up to a maximum of 4.8GT/s, to improve overall system balance and scalability.
- ***Integrated Memory Controller*** – to increase fault tolerance and reduce system downtime and improve system reliability, and provide maximum performance for memory-intensive applications.
- ***AMD-P Power Management*** technologies – with ***Enhanced AMD PowerNow!*** Technology to reduce data center energy consumption, ***Dual Dynamic Power Management*** to reduce idle power consumption and permit per-processor power management in multi-socket designs, ***AMD CoolCore*** technology to cut off power to unused transistor areas to reduce power consumption and lower heat generation, ***AMD PowerCap Manager*** to enable the I.T. staff to put a cap on the P-state level of the cores via the BIOS, and ***AMD SmartFetch*** technology to allow inactive cores to enter a "halt" state and draw less power.
- ***AMD-V technology*** – to enable the translation of virtual to physical addresses in hardware rather than software, I/O virtualization, and the migration of a virtual machine between two physical machines running AMD Opteron technology. This feature is also known as Rapid Virtualization Indexing (RVI).

Source: AMD

specifications as its predecessor. The innovative features in Opteron that enable these improved performance and consolidation capabilities are listed in Exhibit 1, above. However, Istanbul has six cores, 50% more than its predecessor.

- *With six cores*, up to 128GB of RAM and the ability to manage virtual machines dynamically with outstanding performance, Istanbul provides the AMD Virtualization (AMD-V) technology that the enterprise requires to consolidate underutilized systems
- *With six cores*, Istanbul delivers a significant increase in performance over Shanghai, operating within the same power and thermal envelopes, providing outstanding performance per watt results with no additional data center energy consumption.
- *With six cores*, the data center can build a scalable infrastructure that will enable the I.T. staff to utilize platforms fully that support 2, 4, or 8 processors with up to 48 cores and 96 threads, for the growing application needs of the enterprise.
- *With six cores*, Istanbul can replace older platforms and maintain compatibility with existing I.T. infrastructure, running mission- and business-critical applications with increased performance, but maintaining the thermal envelope.
- *With six cores*, AMD has designed Istanbul in two distinct families: the *AMD Opteron 2400 Series* and the *AMD Opteron 8400 Series*. The data center can deploy the Opteron 2400, replacing the Opteron 2300, in two-socket servers, while deploying the Opteron 8400, replacing the Opteron 8300, in four-socket or eight-socket systems. The 2400 processors come with a frequency of 2.2, 2.4, or 2.6GHz for low-power, standard-power, and performance-optimized requirements.

The IBM System x3755 Solution

Once an enterprise has made the decision to replace its current architecture, or expand it, it must decide with whom to partner. With an open systems architecture, there is parity among any number of vendors in terms of application compatibility, which is why it is called "open". The key is to determine what differentiates the competing manufacturers and how those differentiators will affect your TCO. In the case of IBM and their System x3755, the differentiators are compelling, not only in terms of technology, but also in terms of savings to the enterprise on the bottom line.

The System x3755 is a 4U, rack-mounted enterprise server with up to four six-core AMD Opteron 8400[1] processors running at up to 2.6 GHz, at only 75W each. The key words here are "up to", because unlike any other vendor supplying a four-socket server, IBM provides configuration with two, *three*, or four processors. An entry-level dual-processor x3755, running at 2.4 GHz, is web-priced at $10,369. A dual-processor x3755 running at 2.6GHz carries a web-price of $11,439 and is upgradable to three or four CPUs, providing the enterprise with investment protection as business conditions evolve, along with outstanding price/performance to lower data center TCO. IBM's *CPU Pass Thru Card* (see

[1] IBM is offering the x3755 with an AMD Opteron 8431 at 2.4GHz and an Opteron 8435 at 2.6GHz.

Exhibit 2 – The IBM Differentiation

The **CPU Pass Thru Card** enables faster access to information by:

- Replacing the fourth processor, re-enabling the ring topology;
- Decreasing the latency and increasing the memory throughput; and
- Restricting the need of snoop requests to travel only one hop.

Xcelerated Memory Technology counteracts the noise created by read/write signals on the memory bus by:

- Varying the impedance to create a counteracting noise;
- Creating a ghost stub to cancel reflections;
- Maintains the timing and electrical integrity of the signal; and
- Maintains the bus speed at 667MHz for all eight DIMMs on each CPU card.

Source: IBM

Exhibit 2, aboveCopyright (c) 2009 by The Clipper Group, Inc. Reproduction prohibited without advance written permission. All rights reserved.) is the innovative technology that enables the x3755 to grow in increments of one CPU. IBM inserts the card in lieu of a fourth CPU, for near-linear scalable performance and flexibility, resulting in decreased latency and increased memory throughput. These technologies enable the data center to reduce TCO and to have both high capacity and high performance, and not be forced to choose between, but more on that later.

The x3755 supports up to 128GB of DDR II memory, with *chipkill*[2] protection, running at 667MHz. Furthermore, IBM's unique *Xcelerated Memory Technology* (see Exhibit 2), ensures that *all* of the memory will run at 667 MHz, even when the system is fully configured with eight memory chips per processor card. Competitive Istanbul systems typically must clock memory back to 533MHz when the system is configured fully with 128GB. The x3755 also comes with 128 KB of L1 cache, 512KB of L2 cache, and 6MB of L3 cache. The x3755 has a low power consumption rate, utilizing AMD's PowerNow! technology along with IBM's *Active Energy Manager*. IBM's Active Energy Manager measures, monitors, and manages the energy components built into the x3755, enabling a cross-platform management solution. Active Energy Manager extends the scope of energy management to include facility providers to enable a more complete view of energy consumption within the datacenter.

The x3755 is also configured with four hot-swap disk bays, supporting up to 1.2TB of SAS storage. The system includes integrated RAID support for RAID-0, -1, and -10, saving one expansion slot for other use[3]. Additional RAS features include hot-swap and redundant power supplies and fans to help increase system and application uptime, eliminating these elements as a single-point-of-failure. The System x3755 also features pull-down light path diagnostics with IBM *Predictive Failure Analysis* standard, for quick and easy identification of component failures, such as hard disk drives, memory, and CPU, along with fans, voltage regulator modules, and power supplies.

Many data centers are forced to deploy quad-processor configurations in order to achieve the performance required to meet guaranteed SLAs when a dual-processor system simply does not have the required throughput. Obviously, a quad-processor system carries a significant cost burden, especially if the enterprise does not need twice as much performance. However, because of IBM's CPU Pass Thru Card, the x3755 can be deployed with three processors, saving the data center on acquisition cost and operating cost (in terms of software licensing and energy consumption). How does a three- processor x3755 (18 cores) compare to a quad-processor system based upon Intel's six-core Dunnington CPU (24 cores)?

IBM has published results for both a three-processor and four-processor x3755 for the SPECjbb benchmark. The three-processor x3755 running at 2.6GHz delivered 702,512 BOPS[4] and 234,171 BOPS/JVM. Not only is this a 28% improvement over the score of 546,564 BOPS for the previous generation of three-processor x3755, but it also provides the data center with 12%

[2] Chipkill is an IBM technology that protects computer memory systems from any single memory chip failure as well as multi-bit errors from any portion of a single memory chip.

[3] RAID-5 is available as an option.

[4] Business Operation Per Second.

better performance than a four-processor server configured with the Intel Dunnington processor. Comparing a four-processor x3755 is also impressive, with that system checking in with a BOPS rating of 940,674 and 235,169 BOPS per JVM, compared to a competitive Istanbul system with a 937,207 result, probably because of IBM's superior memory speed.

The cost savings for deployment of a three-processor system become even more compelling when the enterprise looks at the software cost for some of the more pervasive enterprise applications. Take IBM's *WebSphere Application Server*, for example, an application licensed by processor count. When a data center requires more performance than is available from a dual-processor server, it can deploy an x3755 with three processors and *WebSphere Business Monitor and Process Server*, and can save $614,000 when compared to the licensing costs for a four- processor server. The opportunity to defer that cost to when even more processing power (the fourth processor) is really needed, with built-in investment protection, is a compelling analysis. Similar software savings are available with *Oracle DB Enterprise Edition* ($285,000), and with VMware's *vSphere* ($3,495).

The x3755 is also a winner when viewed in terms of power consumption. Although not yet listed on the SPEC website, IBM is reporting that the x3755 has the best results in the four-socket category for the SPECpower_ssj2008[5] benchmark. On July 14th, IBM issued a press release announcing a Performance to Power Ratio of 1,165 overall ssj_ops/watt, the highest four-socket score achieved to date.

Conclusion

The IBM System x3755 has been optimized to reduce complexity, improve business performance, enable the enterprise to build a competitive advantage, and drive business growth. The bottom line, quite simply, is that data centers can reduce their total cost of ownership by replacing outdated servers with a new IBM System x3755 with the latest Opteron technology. IBM can provide a growth path via a three- or four-processor system that provides near-linear scalability, enabling the data center to optimize on price and maximize for performance, *at the same time*.

IBM has taken an "open" architecture, where everyone starts on an equal playing field, and added differentiation via technological innovation to distinguish the System x3755 from any other Opteron-based server. By implementing a three-processor model for the x3755, IBM has provided a scalable growth path for the data center while enabling the enterprise to reduce the TCO of the IT infrastructure. If your data center needs to consolidate and virtualize an environment characterized by server sprawl and out-of-control complexity, the x3755 may be the solution that you have been seeking.

[5] SPECpower_ssj2008 is the first industry-standard SPEC benchmark that evaluates the power and performance characteristics of volume-server-class computers.

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.