



IBM x86 Processor Infrastructure — Creating an Enterprise-Class x86 Server

Analyst: David Reine

Management Summary

Ever since the introduction of the Intel x86 microprocessor into the data center, the enterprise has seen the proliferation of the commodity server, from any number of vendors. Software houses around the globe adopted the x86 server as their primary development vehicle, migrating their applications from a variety of proprietary platforms to a heterogeneous open systems architecture. You could choose any x86 server, whether based on Intel or AMD, knowing that their application set would run. In order to gain a foothold in the data center, server vendors have had to remain open within the server, but establish their differentiation in what surrounds the server. That differentiator could be in performance delivered, using different chip sets around the commodity processor to gain in edge in communication on the motherboard; in the surrounding hardware, such as a superior disk array or tape library in the peripheral arena; or perhaps better systems management software, to make it easier monitor and configure. System vendors everywhere were looking for that added-value that would distinguish their platform from just another white box with an Intel CPU.

As the x86 processor has matured, both Intel's *Xeon* and AMD's *Opteron* have added multi-core designs, multi-threaded environments, virtualization features, and energy controls to the basic processor functionality. They also included memory management and intra-system communications capabilities (i.e., between CPUs in the same system), removing the necessity of the external chip set to perform these standard functions. The introduction of multiple cores and virtualization, however, led to another significant advancement within the data center architecture: they enabled the data center to simplify a complex infrastructure through consolidation. **Unfortunately, all of these enhancements have led to another set of challenges, in terms of the memory and TCO issues required to be able to take advantage of the number of virtual machines (VMs) that are available on a new multi-socket platform, even those with only two x86 processors.** In order to achieve a more efficient utilization of the server platform, with more and larger virtual machines, and even larger databases, the processor needs access to more memory, a lot more, and possess even more I/O throughput than ever before. It needs a degree of scalability unheard of before now. It needs to lower energy and management costs and it needs the flexibility to optimize performance for the required workloads. Moreover, the data center needs to do all of this today!

With their fifth version of their CPU-surround architecture, IBM has chosen a path to surround the CPU(s) with breakthrough performance, maximum server flexibility, and simplified management. With their *Enterprise X-Architecture (eX5)* architecture, IBM has managed to develop systems with sufficient resources and innovation to maximize value to the data center, providing a faster integrated data access, maximum flexibility with node partitioning, and optimized configurations for target workloads for quad-socket systems. To learn more about how eX5 can improve your data center performance and lower your TCO, please read on.

IN THIS ISSUE

- Challenges of Enterprise Workloads2
- Enterprise X-Architecture3
- System x Generation 5 Servers4
- Conclusion.....5

Challenges of Enterprise Workloads

The Continuing Need for Increased Performance

The typical enterprise-class data center faces a never-ending set of challenges as it tries to cope with the demands presented by high-performance workloads requiring an increasing store of resources. These workloads are focused on the three-headed monster of improving database performance, increasing transaction processing, and simplifying the IT infrastructure. Furthermore, these workloads must all be capable of being balanced – dynamically – devoting IT resources to the application that needs them the most, no small challenge. Plus, the balancing must be done automatically, in order to reduce the complexity of managing the growing infrastructure.

An End to Wasteful Server Practices

The challenge of improving resource utilization continues to be a key factor in the battle of the IT budget wars. Current x86 legacy servers with a non-virtualized architecture waste a significant percentage of their CPU resources, over 80% in many cases, not to mention the waste of energy required to run the IT environment and cool the data center. The accumulated server sprawl also contributes to an increase in complexity throughout the enterprise, not just in the data center.

Enough Memory to Keep Up with the Processor Performance

Fortunately, multi-core processors, with multiple threads per core, and integrated virtualization features, enable the data center to consolidate multiple applications onto a single multi-socket platform. Unfortunately, many of these newer open systems platforms do not possess the memory capacity required to support (adequately) the dozens of applications that might run effectively in this virtual machine (VM) architecture, especially with requirements for larger VM partitions appearing on the horizon. A lack of adequately matched resources will continue the waste from practices that leave “too much headroom”, especially underutilized processing power. In addition, without adequate memory resources, the data center will continue to struggle with growing databases demanding higher performance while, at the same time, handling increased transaction traffic.

An End to Wasteful Storage Practices

The data center also faces a significant waste

of resources for disk storage. Complex data analysis requires high levels of database performance, which requires very high levels of IOPS¹ to support the data required by demanding applications. Traditional configurations require an increasing number of under-utilized HDDs to deliver the performance required. Very large databases can require literally racks of disk devices to deliver the necessary performance, even if it means seriously over-provisioning the capacity and raising the TCO. Configuration, set up, tuning, and maintenance can add significantly to the TCO of the IT infrastructure. As the number of databases continues to grow, the requirement for additional IOPS increases. This growth can become unsustainable as additional racks of HDDs propagate through the enterprise. In order to reduce the TCO of the storage infrastructure, the data center must replace this escalating drain on energy and the IT budget with solid state drives² to replace the HDDs to improve performance and lower operational costs.

An End to Over-Licensing Software

Processor scalability is another challenge for the IT staff. With total cost of ownership factors contributing to a rising budget deficit, the data center needs to get a handle on escalating software licensing fees and skyrocketing energy costs, both of which affect the infrastructure budget more than acquisition costs. Acquisition costs, however, can also be reduced with the availability of open systems platforms possessing more than two sockets. Quad-socket servers can contribute to budget savings with reduced software licensing and a reduction in the number of platforms deployed in the data center, limiting qualification costs, reducing complexity and energy consumption, and enabling the data center to make better use of existing racks and floor space. This may also help the data center with limited expansion room to defer the move to a new facility, along with the accompanying capital expenses. However, achieving full utilization of a quad-socket server requires significantly more IT resources than commodity two-socket servers have available today.

A Reduction in Infrastructure Complexity

The enterprise data center with open systems server sprawl also faces the task of simplifying the

¹ I/Os Per Second.

² See the issue of *Clipper Notes* dated February 10, 2009, entitled *A New Tier of Storage Appears – Faster, Solid-State Drives State Their Case*, and available at <http://www.clipper.com/research/TCG2009006.pdf>.

IT infrastructure in order to reduce, or, hopefully, eliminate an increasing complexity that builds with every new server that is added to the environment. This complexity hinders the ability of the IT staff to move applications from development, through deployment, and into production. It hinders the capability of the staff to optimize performance for specific workload needs. It increases the TCO as you must perform these tasks manually. The data center needs the flexibility to assign IT resources to the server, to the VM, to the application, that requires them, dynamically.

For each challenge there is an opportunity. In this case, the number of challenges facing the data center requires a flexible, enterprise-class architecture built upon innovative technology.

Enterprise X-Architecture

One company with a solid track record for the consistent delivery of innovation is IBM. IBM has designed their *System x* open systems servers with a unique, internal processor-surround architecture for the past decade, with the first generation appearing in 2001. **IBM is now delivering high-performance x86 servers with a fifth-generation architecture, enabling breakthrough performance, maximum flexibility, and simplified management so that the data center can get the maximum value from both their IT assets and staff, preserving budget.** This new generation of platforms consists of two new four-socket servers for rack-mount and a new blade.

- The entry IBM *System x3690 X5*
- The high-performance IBM *System x3850 X5*
- The new IBM *BladeCenter HX5*

The new version of the *Enterprise X-Architecture* is referred to as *eX5*.

With workloads growing every day and data exploding, the IT infrastructure is being asked to do more and more. 64-bit applications, virtualization, and greater server utilization are accelerating the need for more memory and faster I/O processing every day. Memory-intensive applications are bogging the system down even though there is plenty of raw processing power available. *eX5* has a number of features to boost performance and minimize licensing costs by decoupling memory from the processor, enabling the enterprise to add memory as needed without having to add more processors. Here are its major features.

MAX5

MAX5 is an unprecedented memory expansion capability. It is an innovative technology that enables a generation 5 System x server to connect to an external 1U MAX5 memory chassis. It enables the System x to decouple server memory from system processors to allow you to optimize server performance with up to 2X the memory capacity.

eXFlash

The *eXFlash* provides maximum internal storage performance for System x Gen 5 servers with solid-state disk technology and a new high-speed controller, replacing the IOPS limiting hard disk drives (HDDs). *eXFlash* delivers the same database performance as previous generations for 97% less storage cost for a 1000-user database. *eXFlash* provides the same level of performance as 800 HDDs, delivering a 30X performance increase for local databases. *eXFlash* consumes about 1% of the energy required by HDDs on a performance per Watt basis. This represents a 10X reduction in energy costs for enterprise configurations and may also represent savings in software licensing.³

FlexNode

The *FlexNode* capability provides unprecedented flexibility to the data center. It enables the data center to transform a single four-socket *eX5* system dynamically into two distinct two-socket systems, and also reverse the process, dynamically. This enables the data center to optimize changing application performance needs. What does this mean for the enterprise?

First, *FlexNode* enables the data center to optimize workloads on a project-by-project basis by provisioning System x Gen 5 servers to run interactive applications during the day on a two-processor system and to run batch jobs at night on a four-CPU system. *FlexNode* also may help the enterprise reduce software licensing costs, enabling the data center to execute partitions in order to optimize software licensing charges for applications that are priced per core, lowering the TCO. In addition, this feature allows you to run multiple

³ There are two aspects to consider, both related to the large amount of current HDDs that remain unallocated as part of the strategy to achieve the needed IOPS, as discussed earlier. Your storage management software licensing charges may be based on the physical capacity being managed and not on what you are actually using. Even when 90% is not allocated, you still may be paying for software on 100% of it. Additionally, some middleware and applications also measure base their charges on HDD capacities.

operating systems on the same server without hypervisor overhead. Furthermore, FlexNode enables the data center to isolate workloads to help improve security and assure quality of service.

System x Generation 5 Servers

IBM has always had the broadest portfolio of flexible and scalable rack and blade systems. Now they are expanding that family to address the specific real-world needs of enterprises by maximizing memory, minimizing the TCO, and simplifying deployment. IBM is creating more efficient and cost-conscious servers with:

- The latest high performance Intel *Nehalem*⁴ processors;
- A lower entry point for enterprise-class processing;
- More memory than currently available with unique extension capabilities;
- Maximum flexibility with FlexNode node partitioning; and
- Configurations optimized for specific target workloads.

With its MAX5 technology, IBM has dedicated more DIMM slots for each of their Generation 5 servers than we have seen in other high-end open systems platforms. The additional memory slots enable you to improve consolidation ratios, deploying more applications per server, lowering the TCO for applications charged on a per core basis.

Furthermore, you can configure up to three eXFlash chasses per X5 server, with up to 1.6TB of data per eXFlash. This represents up to 240,000 IOPS per X5 server. Each eXFlash can be configured with RAID 5 or RAID 6 for redundant data storage, or without RAID for high-throughput requirements. The eXFlash is hot-swappable with a front access to simplify maintenance for System x Generation 5 servers.

Each Generation 5 server also comes with the latest version of IBM *System Director* to provide a single, consolidated point of management with remote access.

IBM System x3690 X5

IBM's System x3690 X5 possesses a new design, offering the best density for enterprise computing at an entry price. It enables the

data center to increase productivity and reliability with an easy-to-deploy and manage platform. The x3690 comes standard as a 2U, 2-socket server with 32 DIMM slots in the entry-level configuration. What distinguishes this server from other entry-level platforms are its levels of scalability and flexibility.

You can upgrade this entry-level server from a 2U, 2-socket architecture to a 4U, 4-socket configuration. Furthermore, you can add a 1U MAX5 memory expansion unit enabling the x3690 with 64 DIMMs in a 2-socket configuration or 128 DIMMs in a 4-socket deployment. Processors and memory are scalable independently, delivering maximum flexibility. You can deploy more and larger VMs, allowing for faster database response without burdening the enterprise budget with the cost of unnecessary processors. You also can choose to install smaller, less expensive DIMMs, if memory cost is the primary issue instead of capacity.

IBM System x3850

IBM's System x3850 X5 provides the enterprise data center with a new level of scalability to meet the challenges presented by new ways of doing business. The x3850 enables the enterprise to manage costs, while meeting the changing requirements of an evolving business. It is a rack system offering the performance, flexibility, and scalability that your enterprise demands with an affordable entry price and the capability to expand with your business. It can grow processing capability from four sockets to eight, with independent scalability for memory and storage.

X5 enhancements (over the Gen 4 version of the x3850) enable IBM to deliver an enterprise-class system with simpler qualification, faster deployment and easier administration than ever before. With FlexNode, you can reconfigure the x3850 dynamically to meet the changing application needs of the data center.

Configurable as either a 4- or 8-socket rack system from the factory, the x3850 X5 has the independent configurability to satisfy the most demanding environments. With 64 DIMM slots standard for a 4-socket system, the x3850 X5 can scale to 96 DIMMs for a 4-socket configuration or 192 DIMMs for an 8-socket system, with seven I/O slots. With the additional memory, the data center can deploy more applications per system, more applications per core, and lower the cost of license fees for those applications that are charged on a per core basis.

⁴ For more on Intel's Nehalem, see [The Clipper Group Navigator](http://www.clipper.com/research/TCG2009017.pdf) entitled *IBM Lowers TCO in the Data Center- Innovation Surrounding Nehalem Scores Big*, dated March 30, 2009, and available at <http://www.clipper.com/research/TCG2009017.pdf>.

IBM HX5 for BladeCenter

IBM has taken the BladeCenter HX5 beyond the bounds of “standard” with unique innovation, extending the value of eX5 to BladeCenter. Built with the same qualities of efficiency, cost-consciousness, and productivity as its rack-based cousins, the HX5 significant value for your business-critical applications. It delivers enterprise-level compute performance in a format that provides maximum memory and I/O bandwidth for the data center.

Configurable as a 2- or 4-socket blade, the HX5 has 16 DIMM slots standard with a 1U MAX5 memory expansion blade to support up to 40 DIMM slots for a 2-socket server or 80 DIMM slots for a 4-socket system. This represents much more memory than has been possible before. The HX5 also has two onboard solid-state drives to improve I/O speed and reliability. In addition, each HX5 blade has two PCIe slots and eight I/O ports for communication flexibility, with *Open Fabric Manager* available to manage the I/O flow.

Conclusion

If the challenge to your enterprise data center is to get the maximum value out of your IT infrastructure budget, then you need to consolidate underutilized infrastructure onto superior virtualized platforms. In order to get the highest utilization rate the data center must maximize the amount of memory available to the latest open systems processors and throw enough storage I/O at the environment to satisfy the most demanding applications. Maximizing memory, minimizing cost, and simplifying deployment is exactly what IBM is doing to optimize their fifth generation eX5 architecture and their generation 5 system x servers.

With X5, IBM is enabling the enterprise data center to reduce their server count in half for a given workload, reduce storage costs by up to 97%, and reduce software license fees by half, lowering the TCO for the IT infrastructure. Using the latest Intel Nehalem processors along with IBM innovation in memory and storage architecture, your enterprise datacenter can achieve tremendous productivity gains through more efficient use of memory and I/O. IBM has enabled the data center with the flexibility and scalability to grow processing power, memory, and storage independently, to ensure that your business-critical applications have the resources they need when they need them, without the requirement to over-provision resources that are not needed.

IBM is changing the way the IT game is being played in the enterprise data center. They are changing the high-performance paradigm, providing dynamic access to IT resources dynamically. If your enterprise is faced with the inefficient deployment of servers, memory, and storage for your HPC applications, look at the fifth generation of IBM System x servers, they may be able to change the way that your enterprise can deploy your next production applications.



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 a Senior Contributing Analyst for The Clipper Group. Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. In 2002, 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.

Disclosures

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.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper's web pages, posting of Clipper documents on other's websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts' travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper's analysts.

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.