THE CLIPPER GROUP
# Navigator™

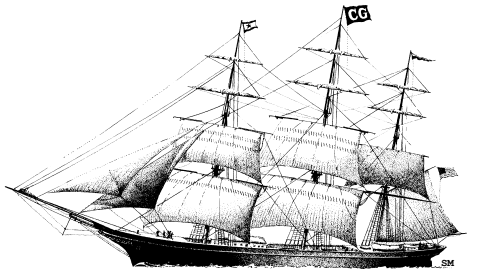

*Navigating Information Technology Horizons*℠

Published Since 1993 Report #TCG2011019 May 20, 2011

# Redefining the High-End Server — IBM Upgrades System x for Large Workloads

Analyst: David Reine

## Management Summary

Whenever you look for a new car, there are a lot of questions to be asked before you step one foot into the dealership, unless of course you are looking for a little, red, convertible sports car that can go from 0 – 60 in less than four seconds, and you don't care about gas mileage. The most important of these questions might be: what kind of workload do I intend to put inside my new auto? (Well, you might not put it in those terms, but bear with me.) If you are looking to carry a family of four, with a dog, you might look one way. If you are looking to carry a lot of cargo, you might look another. If you are concerned about gas mileage, as the price for a gallon of gas exceeds $4.00, and supposedly on its way to $5.00, that too can influence your decision. In some cases you will be able to find a vehicle that can be reconfigured dynamically, adding more passenger seats as your family grows, or removing them, enabling more cargo space. Your car might need to adapt from two-wheel drive to four-wheel drive, if your idea of a family vacation is a week in the desert. Being able to conveniently adapt your ride to a variety of workload requirements may be the decisive selling point in your next car-buying experience.

Being able to adapt to a variety of workload environments is nothing new to the CIO or manager of any business- or mission-critical data center. In a world of Internet commerce, there is no such thing as off-hours. Your IT infrastructure has to be available 24x7x365 in order to ensure that your enterprise website is always accessible. The infrastructure needs to be highly-reliable and highly-performant to enable the IT staff to migrate business- and mission-critical workloads from more expensive proprietary UNIX systems to lower-cost x86 servers. These servers have to be very flexible to ensure that they can handle changing daily and seasonal peak workloads. Memory and processors must be flexible in how they can be allocated and changed, dynamically, to reconfigure systems to adapt to changing workloads, to move virtual machines from one server to another for high-availability or workload balancing, or to be able to conserve energy when conditions warrant. What this means is that the capacity of processing power and memory allocated to partitions must be "dial-able" up and down to meet changing needs. Also, for some data centers, it might be advantageous to increase the server's installed capacities at a later date. This would allow servers to evolve to meet growing needs of applications and users. There are many vendors of commodity servers using standard x86 technology that can satisfy some of these requirements, but only one vendor has integrated innovation into a standard architecture to satisfy all of these requirements at the highest level, and that is IBM.

IBM now is delivering its fifth generation of *System x* servers, adding performance, reliability, and functionality at each stop along the way. This latest announcement has upgraded IBM's System x family with the latest innovations from Intel, that being the 10-core Intel *Xeon processor E7*. With this new System x platform, the data center now has the tools that it needs to consolidate the IT infrastructure with processor virtualization and support for the largest databases. To learn more, please read on.

## IN THIS ISSUE

## Data Center Optimization Requirements

The issues confronting every CIO and data center manager are the same – *trying to do more with less*. The growth in the number of servers proliferating throughout the data center, and the enterprise as a whole, is causing a dramatic effect upon the IT budget for every data center, SMB or enterprise, and that budget tends not to be growing. This is not just caused by the acquisition cost of the server platforms; it also is greatly affected by the total cost of ownership (TCO) of the IT infrastructure as a whole. This includes not just the servers, but also the storage devices, networking, etc., which are consuming vast amounts of administrative support and maintenance in the data center, not to mention energy and floor space. With the cost of energy for our cars racing past $4.00 per gallon, we can expect the cost of the electricity that runs and cools our data centers to soar as well, albeit a little more slowly. Moreover, if the data center runs out of energy, or even floor space, necessitating the construction of a new data center, it will have enormous ramifications, possibly affecting your job. If it is possible, the requirement to consolidate data center resources and optimize the workloads running on more reliable and more secure servers is becoming important, in order to meet data center limits, lower that TCO, and help to solve real business challenges.

Optimizing client workloads is becoming a key factor in the consolidation and virtualization of the enterprise data center. This means providing sufficient resources to meet the needs of the application and users, as their needs change. The goal is to prioritize what is being run without over-allocating resources that might better be applied elsewhere. The IT staff has to raise the limit not only on the number of workloads running on the application servers, but also the size of the virtual machines. The flexibility to change the underlying resources dynamically to more powerful processors, with even more cores and threads, enables the data center to run more and more business- and mission-critical applications, in less space. The number and processing power of the CPUs within the server are important, but the ability to allocate flexibly these resources is more important. The more cores that there are within the processor, the more threads there are to handle more applications (by not having to wait for a thread to become available). More threads mean that the IT staff can run more applications in parallel (i.e., more applications at the same time on the same processor), maximizing the flexibility of the architecture, improving dynamic growth, and leading to higher reliability.

However, having more high-powered processors is not a be-all, end-all in isolation. You also must have a more memory that is more flexible. This will enable the IT staff to not only roll out more applications per core, but it will enable larger virtual machines (VMs) and larger databases to be deployed, without choking. How much memory? How high is high? The same thing is true for critical business applications will take whatever memory you can allocate. The trick is not to allocate so much that the extra increment provides little tangible benefit. Typically, the more memory you have, the more applications you can run.

With flexible processors and memory, the IT staff can not only improve the utilization of the processor, but also improve the utilization of the entire system, CPUs, memory, I/O, and storage. In rack systems, this means more processing power for every "U" in the rack, and, therefore, less floor space. For blade-based systems, it means better density and performance within the blade server and fewer blade chasses required.

This flexibility will also have a significant effect in the cloud, where the sharing of resources is a requirement. Moreover, with improved encryption, the E7 processors can protect the data better, reducing risk of theft.

To do all of this and still lower the TCO of the IT infrastructure, requires an improved open systems architecture, with more powerful, more efficient, and more reliable processors to lower the impact of recurring costs. With the desire to run more virtual workloads that are larger and more critical, x86 servers are taking on more demanding roles, including high-end critical applications. The newest member of the x86 open systems architecture family to directly address these challenges is the Intel *Xeon E7*, known internally as *Westmere EX*.

## The Intel Xeon E7 Processor

Once again Intel has raised the ceiling of processor performance and functionality. In order to meet the demanding roles being thrust upon them, x86 servers are now being designed to satisfy the needs of high-end business applications. To meet that need, Intel has designed the Xeon E7 family of processors to provide the

data center with a top-of-the-line, enterprise-class server to manage and secure crucial business-critical applications, while exceeding the demands of most mission-critical IT environments. The E7 family uses Intel's *Nehalem architecture* as a springboard to deliver the ideal combination of performance and scalability[1] for your most demanding workloads, using a *scale-up architecture*. Scale-up enables the deployment of fewer, more powerful servers with the capability to handle more demanding applications, with built-in headroom to handle peak workloads and normal growth. Today, some enterprises run their mission-critical applications on RISC[2] processors, such as IBM's *POWER*, and Oracle's *SPARC*, or on an EPIC[3] processor, such as HP's *Itanium*. The IT staff will have to weigh the commitment of these vendors to their platforms against the costs of migrating to an open x86 environment with *Linux* against the expected performance gains and costs of remaining on their legacy platforms.

These new x86 servers scale well beyond the capability of older x86 servers, but the actual TCO of the solution depends on more than just the acquisition cost of the platform. The IT staff must consider the total number of servers needed, the consolidation and virtualization density desired, the ultimate scalability (in terms of cores and memory of a single platform), energy costs, software license costs, migration costs, etc. All of that needs to be considered before deciding that just because you can do something on x86, you should.

With improved range of performance, increased memory, and expanded I/O functionality and capacity, your servers now can adapt readily to changing short-term business demands while addressing the requirements of long-term business growth. The E7 family enables the data center to consolidate and virtualize more applications on their most capable servers, where advanced reliability and security features work to maintain data integrity, accelerate encrypted transactions, and maximize the availability of business-critical applications, all within the energy envelope of 130 watts of the previous-generation CPU (*Xeon 5600*, also know as *Westmere EP*). In addition, the multi-core E7, using a 32nm process technology, enables the data center to lower the TCO of the IT infrastructure with up to 95% lower operating costs on four socket servers than traditional single-core CPUs.

In comparison to the Xeon 5600[4], the E7 has outstanding functionality with up to 10 cores per socket, as compared to six cores. Each E7 socket supports up to 20 threads and up to 16 low-voltage DIMMs (up to 4TB for 8 sockets) with 30 MB of shared L3 cache. This compares quite favorably over the Xeon 5600 with up to 6 cores, 288GB of memory, and 12MB of L3 cache. In addition, the E7 has more memory lanes for performance and includes memory buffers, enabling the system to keep the clock speed constant. The E7 also includes a *Memory Check Architecture*, to improve reliability while it retains the Xeon 5600 capabilities in terms of interconnect (QPI), memory type (DDR3), and micro architecture (Nehalem). It also retains many of the same innovations from the 5600: *Intel Turbo Boost*, *Intel Hyper-Threading*, *Intelligent Power*, and *Intel Virtualization*. The E7 improves upon the reliability, expandability, scalability, and I/O capacity of the 5600. In addition, the E7 delivers the same reliability and data security as the Intel *Itanium Processor 9000*.

In terms of RAS, the E7 has advanced features that help to protect both systems and data more effectively, lowering the TCO by reducing the need for malware protection and protecting the enterprise from data breaches and data loss. These include seven new *Advanced Encryption Standard New Instructions* (*AES-NI*), to reduce the performance penalty usually experienced with pervasive encryption, while enabling encryption software to run faster, with better data protection. In addition, the E7 includes the *Intel Trusted Execution Technology* (*Intel TXT*) to increase malware protection through measured launches into "known good states". Intel TXT enforces control of the platform, enabling isolation in the boot process and while complementing runtime protections.

## IBM's System x Solutions

Many data centers have already faced one challenge – consolidation, virtualizing many web-facing applications that were easy to virtualize, i.e., they were the low-hanging fruit. From this, enterprises have realized significant benefits in the form of reductions in capital

[1] The E7 can be configured with 2 to 256 sockets.

[2] Reduced Instruction Set Computing.

[3] Explicitly Parallel Instruction Computing.

[4] Westmere EP.

expenses and operational expenses. Now, the virtualization of business-critical applications and databases is taking center-stage, along with building cloud infrastructure. This is happening as part of an effort to improve the utilization of business-critical servers, while lowering the TCO and still remaining compliant with industry and government standards. For many, these are new challenges, with significant expected benefits in addition to the improved TCO, all of which will affect the quality of service and roll-out time.

IBM has integrated Intel's Xeon E7 into four models of their System x family, with an estimated 20%-30% improvement in performance. The *System x3690 X5*, the *System x3850 X5*, the *System x3950 X5*, and the IBM *BladeCenter HX5* blade server. However, IBM did not just take the E7 and stick it into an everyday commodity server system. IBM has surrounded the E7 with the fifth generation of their *X-Architecture*, *eX5*, to improve overall performance, scalability, reliability, and security. IBM X5 servers with the E7 processor are designed to support the data center today, and for years to come, protecting the investment being made today. What are some of the features that constitute IBM's eX5 CPU-surround strategy?

IBM has implemented the following features into their System x servers.

- ***MAX5*** is the memory access feature for eX5 that enables the IT staff to deploy an additional 32 DIMM slots for the x3850 X5, x3950 X5, and x3690 X5, for a total of 64 DIMM slots, and a total of 40 DIMM slots for the HX5. This provides the x3x50 servers with the capability to scale up to 3TB and the HX5 blade up to 640GB. This provides the X5 servers with the memory capacity to maximize VM density, increasing the scalability, utilization, and performance of these platforms, while potentially lowering licensing fees. MAX5 also enables the data center to deploy memory with lower cost DIMMs, while retaining the same capacity.
- ***eX5 eXFlash*** can be configured with up to eight solid state drives (SSDs) per eXFlash, including new 50 and 200GB SSDs, to maximize performance and storage, replacing traditional spinning media (hard disks) with ultra-fast SSDs. An IBM System x3850 X5 can support up to two eXFlash packs, or three in a System x3690 X5. IBM's SSDs run at 240K IOPS[5] with up to 97% lower cost and up to 99% lower power than the HDDs that they would be replacing to achieve the same level of performance in high-intensity I/O configurations, with nine times faster response time. With no moving parts, the SSDs also add significantly to the reliability of the system, lowering the TCO.
- ***FlexNode Partitioning (eXA Scaling)*** enables the IT staff to deploy an eXA-scaled system as two independent systems, or as a single system, without removing any cables. It is available today on the HX5 blade and will be available for the System x3850 dual node system with the MAX5 option. It facilitates flexibility in configuration, without operator intervention, in changing from a scale-up application to a scale-out environment and back. The partitioning can be done remotely through the integrated management module.
- ***Improved RAS characteristics*** are provided by *Double Data Device Correction* (double chipkill), *Memory ProteXion* (which improves the reliability up to four times the memory protection of an industry standard server with off-the-shelf DIMMs), and an Advanced Encryption Standard (AES) that allows faster encryption by applications.

This fifth generation of IBM System x servers are designed for "big data" to keep IBM customers ahead of the curve. These IBM innovations enable businesses to convert new streams of data into new revenue, with the knowledge that these servers can still be in service three years down the road, protecting the investment being made today.

### *IBM System x3690 X5*

The IBM System x3690 is ideal for the data center looking to get four processors worth of performance, based upon the performance of the previous generation, in a two-processor server, delivering outstanding performance and reliability with scalable memory and storage, without increasing TCO, space, or energy, increasing the performance per watt. The x3690 combines the latest in Intel processor technology with IBM's fifth-generation X-Architecture (eX5). The x3690 X5 is configurable up to 30MB of cache, up to 20 cores, running at 2.4GHz, and 32 DIMM slots within the system, or 64 DIMM slots using MAX5 functionality, doubling the number of VMs supported, for a total memory

[5] I/Os per second.

capacity of 2TB. Reducing CPU count enables the data center to cut software licensing fees, likely reducing TCO further.

It has five PCIe slots and room for up to 16 2.5" HDDs (up to 9.6TB per chassis) or 24 1.8" SSDs (up to 4.8TB). The x3690 comes with two integrated GbE ports and an Emulex 10GbE adapter. It has IBM *OnForever* reliability features, such as hot-swap power supplies, fans and HDDs, advanced light path diagnostics, and *Predictive Failure Analysis*, to help your users avoid interruptions on mission-critical workloads. RAID-0 and RAID-1 are standard, with RAID-5 and RAID-6 available as an option. The x3690 also comes with IBM *System Director* to simplify management of both physical and virtual resources.

The x3690 can be pre-configured as a *Workload Optimized Solution* for faster deployment and, therefore, faster time to benefit. The x3690, with eX5, can also be configured to handle database workloads or virtualized environments, better than an industry-standard server, based upon such features as MAX5, which can expand memory to enable more VMs. The x3690 X5 has an entry price of $5,900.

### *IBM System x3850 X5 & System x3950 X5*

The IBM System x3850 X5 and the x3950 X5 provide the server reliability that any data center may need to maximize system memory, minimize cost, and simplify the deployment of mission-critical application workloads. A modular design facilitates customizing the x3850 and x3950 for current requirements while retaining the flexibility to scale larger as workloads increase. The System x3850 and x3950 can be pre-configured and optimized for a variety of compute-intensive workloads, including consolidation, virtualization, and enterprise database applications. In fact, IBM has pre-configured one version of the x3950 X5 specifically for *SAP High Performance Analytics.*

Like the x3690 X5, the x3850 X5 and the x3950 X5 are based upon Intel's new E7 architecture and IBM's eX5 server-surround environment. They, too, are flexible and reliable servers, with IBM's OnForever features (described above), along with the same RAID features of the x3690, high performance, a low TCO, and an energy-friendly architecture. The E7 comes with six, eight, and ten-core options at up to 2.4GHz. Scalable memory enables the IT staff to manage memory-intensive applications with higher utilization and greater throughput. This enables the data center to use these servers as a tool for consolidation, without worrying about how many "eggs" are in this "one basket". They, too, have two integrated GbE ports and an Emulex 10GbE adapter.

Unlike the x3690, these systems have a 4U chassis format[6]. They come as two-socket servers, but can expand to accommodate four processors, or even economically scale to eight processors by adding IBM's *FlexNode* technology for optional two-node support. With MAX5, either system can scale memory up to 192 DIMM slots, supporting up to 6TB. This offers failover and partitioning flexibility, enabling the data center to dynamically unite and separate dual node servers with MAX5. Combining this expanded memory capability with an embedded *VMware vSphere Hypervisor* enables the IT staff to increase the size and number of VMs along with the size and speed of databases, improving response time and meeting SLA expectations.

Both systems come with seven half-length PCIe slots and have 8 hot-swap disk bays for 2.5" SAS drives, or 16 hot-swap SAS SSDs, for a maximum HDD capability of 4.8TB per chassis, or 3.2TB of SSDs.

### *IBM BladeCenter HX5*

IBM has designed the BladeCenter HX5 to combine the benefits of BladeCenter and the IBM X-Architecture, with the new Intel E7 technology, to provide the data center with workload optimized systems for both virtualization and databases. The HX5 delivers outstanding performance and utilization in a blade form factor for those environments.

The HX5 is a two-socket blade server, scalable to four, supporting up to 40 cores at up to 2.67 GHz, using the Xeon E7-2800 and Xeon E7-4800 CPUs. The HX5 supports up to 40 DIMM slots, when combined with MAX5, supporting up to 640GB of memory. It has maximum flexibility with room for one standard PCIe and one high-speed PCIe daughter card, for a total of eight ports of I/O to each blade, including four ports for high-speed I/O. This is scalable to 16 ports in a four-socket, double-wide form factor. The single-wide HX5 has two non-hot-swap bays for 50GB SSDs, with optional RAID support for RAID-0, RAID-1, and RAID-

[6] The System x3x50 requires 10U when deployed with two nodes and Max5 functionality.

1E[7]. The HX5 comes with IBM's integrated systems management processor.

The IBM X5 family of servers have performed very well in industry standard benchmarks. For example, X3850 X5 has the highest result for a non-clustered server in the *TPC-H 1000GB* benchmark, resulting in both outstanding performance and price/performance as a result of the integration of E7 with IBM's eX5 architecture. The x3850 X5 has also achieved a world-record four-processor result on the two-tier *SAP SD* standard application benchmark, 14,000 SAP SD users.

IBM has also announced the IBM Systems x solution for *SAP In-Memory Appliance*, *SAP HANA*. Certified by SAP, the systems enable SAP analytics of SAP ERP data on the fly with proven capabilities of processing 10,000 queries per hour against 1.3TB of data. IBM and SAP cooperated closely to implement the advanced replication scenario for *DB2*-based customers. DB2 is SAP HANA-ready and can efficiently replicate data into SAP HANA in near real time using the Sybase Replication Server. Customers with ERP systems on DB2 can support demanding business applications seamlessly for real-time reporting based on the latest available data with little administrative effort. These benchmarks reflect very impressively on IBM and what can be done with a commodity x86 microprocessor, when surrounded by an innovative architecture.

## Conclusion

Is your enterprise data center approaching crisis mode? Do you have floor space for the IT infrastructure that your IT staff has to add immediately? Is there enough energy coming through the walls of your data center to drive that infrastructure and cool the data center environment? If the answer to the first question is "yes", and the rest are "no", then you need to investigate a further consolidation and virtualization of your enterprise infrastructure.

How can you do this and stay within budget, not only for the acquisition of the new servers, but also for the TCO of this new infrastructure? As discussed, IBM has a variety of the highest performing x86 servers. These include both rack-mounted chassis and server blades, at a variety of configurations and performance levels to meet your application performance requirements. However, IBM's high-end x86 platforms, while meeting all the standards requirements for Microsoft *Windows*, *Linux,* and VMware's solutions, have the unique X-Architecture to differentiate themselves from their competition.

With a value-priced family of servers with superior performance, IBM distinguishes itself from the competition with innovation and the right server for any of your workload levels. All you need to do is to determine your workload requirements and budgetary limitations and choose the system that best fits into your data center environment. With the X5 servers from IBM, along with the E7 processor from Intel, you can position your enterprise with the right server for your workload level today, knowing that you have the scalability to meet the workload levels of tomorrow. If any of this rings a bell for you, check out IBM's eX5 family of servers.

[7] RAID 1E uses 2-way mirroring on an arbitrary number of drives.

### 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.