Research Report

A Major Change in x86 Server Design: IBM X6 Platforms

Introduction

To date, one of the biggest shortcomings in x86 system designs has been lack of memory. Pre-Intel EX systems have generally topped-out at around the 2-3TB of main memory range – forcing x86 servers to constantly swap data between memory and storage when working with large databases. Furthermore, these servers have had to use the PCIe bus to read/write to/from disk (instead of using a memory channel to access data in cache). This combination of constant data swapping and the use of the PCIe bus architecture has resulted in input/output (I/O) bottlenecks and has slowed I/O per second (IOPS) performance. Thus, these I/O bottlenecks have forced enterprises to buy more servers to get work done – and, consequently, to spend a lot more money on software licenses (because software is priced per core, so more machines = more cores = higher license costs).

Buyers of x86 servers also have faced application/database headroom challenges. These challenges have precluded servers based on pre-Intel Xeon processor E7-8800/4800/2800 v2 technology from running large in-memory databases; from serving large-scale business applications; from supporting a wide variety of memory- and storage-intensive business analytics applications; and from supporting larger virtual machine configurations.

With the introduction of the new class of Intel Xeon architecture (referenced above), Intel has made it possible to address memory configurations greater than 3TB (IBM's new X6 servers now address up to 12 TB of memory/flash cache) – opening the door for enterprises to use x86 architecture for new larger memory workloads including: 1) large-scale ERP (enterprise resource planning); 2) large database business analytics; 3) large database serving, and, 4) high-IOPs virtualization.

> ***Now that x86 Xeon-based E7-8800/4800/2800 v2 and beyond servers can finally address larger memory configurations, we are expecting to see new generations of large memory x86 servers come to market. And, as could be expected given its constant leadership position in 4+ socket-servers IBM is one of the first major vendors to introduce servers based on this new generation of Intel Xeon processors. The new IBM X6 family of servers are based on IBM's X6 architecture and are offered in both rack and a blade-based platforms***

Enterprises that adopt these new large memory rack/Flex System configurations can expect to:

- Lower computing costs because fewer servers will be needed to do the same amount of work (due to increased processing power and the elimination of I/O bottlenecks);
- Reduce data center server sprawl due to the need for fewer servers – and because significantly more virtual machines can be deployed on large memory X6 designs;
- Host new memory-intensive large scale applications on x86 architecture.

These enterprises can also expect to see improvements in service delivery (due to better performance), as well as simplified management (due to new automated management facilities introduced by IBM that are available with its X6 platforms).

In this *Research Report*, *Clabby Analytics* describes*:*

- The market situation with respect to today's x86 server limitations;
- How IBM X6 offerings help address these limitations;
- How we expect X6 servers to be used (use cases); and,
- The benefits that we expect this architecture and associated IBM value-added management software will deliver to users of X6 systems.

Market Positioning

In the past, three obstacles have blocked x86-based servers from competing more favorably at the top of the server market with higher-end, more vertically scalable reduced instruction set computing (RISC) such as Oracle's UltraSPARC and explicitly parallel instruction computing (EPIC) servers such as Intel's Itanium. These obstacles have been:

1. The *amount of processing power* an Intel x86 processor could deliver;
2. The *amount of data* that could be addressed in-memory; and,
3. The *speed at which data could be served* to the CPU through the I/O subsystem.

Any one – or all of these – created bottlenecks that slowed high-end x86-based server performance as compared to RISC and EPIC systems.

Between 2006 and 2009, Intel started to address two of these shortcomings by: 1) introducing 64-bit architecture – thus increasing the amount of memory that could be addressed from 4 gigabytes (the 32-bit maximum) to a theoretical maximum for 64-bit architecture of 144 petabytes; and, 2) by implementing multi-core processors to increase processing power (starting in 2009 with the arrival of the Nehalem-class Xeon processors and continuing today with consistently more and more powerful multi-core processor implementations).

With the increase in addressable memory, it became possible to place a lot more data in memory and cache where it could be processed more quickly by increasingly more powerful CPUs. *But, to date, due to pre-Intel Xeon processor E7-8800/4800/2800 v2 technology design limitations, most industry standard x86 server implementations maxed-out at between 2 and 3TB of main-memory*. (Note: The previous generation of IBM X-Architecture (eX5) used MAX5 architecture to deliver an additional 32 DIMMs – thus increasing the amount of memory available to current generation IBM System x3850 X5 and BladeCenter HX5 servers well beyond the 3TB limitation). Still, the amount of memory that x86 servers could support has remained an impediment for x86 server owners who have wanted to deploy more scalable, memory-intensive, high IOPS applications on x86 systems. That is – until now – with the introduction of the new IBM X6 server offerings.

About the IBM X6 Server Portfolio

For over a decade IBM has differentiated its high value x86-based servers using its Enterprise X-Architecture (EXA) to help scale the size of x86-based server configurations, to extend memory, to improve input/output per second (IOPS) performance, to improve virtual server performance, and to improve system resiliency. *Enterprise X-Architecture is the reason why IBM System x servers have consistently claimed the top spot in the 4+-socket segment of the x86 marketplace*.

In early 2014, IBM introduced the sixth generation of EXA (the X6 family consisting of System x 3850 X6 and x3950 X6 servers [4- and 8-socket servers]. In May, 2014 IBM introduced its new Flex System x280 X6, x480 X6 and x880 X6 compute nodes (in 2-, 4- and 8-socket configurations). The X6 servers feature newly redesigned modular 4-socket and 8-socket rack servers that offer differentiated design flexibility – as well as the Flex System X6 Compute Node

portfolio featuring 2-, 4- and 8-socket models (IBM's rack and Flex System nodes are illustrated in Figure 1).

Figure 1- IBM's X6 Server Portfolio

Source: IBM Corporation – May, 2014

IBM also announced that it had greatly expanded the amount of data that can be held in cache using IBM's eXFlash memory-channel storage. (Note: By automatically placing more data into cache memory, I/O bottleneck problems suddenly start to disappear. And more data in-memory cache leads to significantly faster IOPS performance and produces faster computing results). IBM's new X6 platforms enable terabytes of data to be cached into Flash storage (technically block storage) -- processors can then exploit large amounts of data in cache using fast memory channel connections. It is IBM's intention to offer eXFlash memory-channel storage for the new IBM Flex System X6 Compute Node Portfolio. This new offering will enable clients to increase database performance, realize higher levels of storage consolidation, and further scale virtualization on a single server.

> ***The big news in this X6 systems announcement is the introduction of the new IBM eXFlash memory-channel storage (dual inline memory modules) that can be directly attached to the memory bus to significantly improve input/output per second (IOPS) performance. New IBM X6 servers enable customers to plug up to 12 TB of eXFlash memory-channel storage directly into memory slots. We believe that this* eXFlash memory-channel storage *has the potential to double database performance (as compared to accessing external data stored on solid state or hard disks using an input/output bus such as the PCIe bus).***

In addition to introducing a new generation of Enterprise X-Architecture, IBM has also reengineered its rack and blade-based server designs. New IBM rack servers can support 96 DIMMs using 64GB LRDIMMs or 6TB of memory in 4-way configurations. On 8-way rack platforms, IBM X6 architecture can support up to 12TB of memory. In addition, the new "book" design of the rack servers allows buyers to mix and match compute, storage and network (I/O) books to the requirements of given workloads, providing buyers with a lot of flexibility in systems designs.

> ***By allowing eXFlash DIMMs to be plugged into memory slots, the IBM X6 design greatly improves data read/write speed by using Flash as opposed to hard disk – helping to eliminate I/O bottlenecks. We expect the use of Flash DIMM technology to result in tremendously faster IOPS performance as compared with other competitor's x86 implementations.***

With more than three times the amount of memory of other competing x86-based solutions, memory-intensive applications can now reside on x86 architecture – IBM on X6 platforms. Combining these speed improvements with significantly more memory means that x86 architecture can host enterprise workloads that require high availability and significantly increased memory!

How We Expect X6 Platforms to Be Used (Use Cases)
IBM is focusing its X6 server marketing and sales efforts on four market segments: 1) large scale business applications (such as run-the-business enterprise resource planning applications); 2) in-memory or high memory business analytics application environments; 3) high performance databases; and, 4) environments that require access to large populations of virtual machines (such as some cloud environments).

Large Scale ERP
By providing x86 processors with access to more memory using IBM X6 platforms and eXFlash memory-channel storage, IBM X6 servers will be able to scale higher and perform better when running business applications such as SAP Business Suite. Although benchmark data was not available at the time that this report was published, we fully expect that thanks to larger memory configurations, IBM will soon be able to claim that its X6 servers are the performance and scalability leaders when running ERP applications such as SAP Business Suite.

Analytics
Over the last several years IBM has become a leader in bringing large memory x86 configurations to market. With the ability to address over 12 TB of data in its new X6 architecture configurations, IBM is well positioned to make the claim that it offers the best x86-based server solution for the SAP in-memory analytics environment (known as HANA).

Also, thanks to patterns (streamlined configuration/deployment practices) as well as other integration work, deploying large-scale analytics applications should be far simpler on large memory, large-scale X6 systems as opposed to typical x86-based servers. For instance, using configuration patterns, IBM has enabled users to significantly reduce the time it takes to deploy validated Microsoft SQL Data Warehouse implementations. We expect numerous other patterns to be developed that will streamline the deployment of other business intelligence/business analytics offerings on X6 servers.

Database Performance
When comparing IBM Flex System x280 X6 compute node to IBM's HX5 blade offering, IBM's Flex System large memory configuration offers 3 times the memory capacity and 50% more compute resources. With additional memory and computing power IBM's X6 architecture has been able to achieve leadership database performance, including:

- The number one position in the TPC-E transaction processing benchmark (5576.27 tpsE @ $188.69/tpsE; ; and,
- The number one position in the SAP SD 2-Tier benchmark at 25,000 users with .98 second response time (and 136, 670 SAPs delivering 2.7 million business processed line items per hour).

Virtualization
In virtualization, IBM can now claim the number one position in virtualization performance – as demonstrated by the SPECvirt_sc2013 benchmark (2081 @ 116 VMs). This is because large data caches make it possible to support for substantially more virtual machines on X6 architectures.

And more virtual machines on fewer cores should help IT executives significantly decrease their virtualization infrastructure licensing costs because X6 architecture enables more work to be performed on fewer cores.

> ***The ability to process work faster has an interesting side effect: there is the potential to decrease solution deployment costs. (These savings come from reduced software licensing costs and lower capital expenditure). A big reason why software licenses and capital expenditure costs decrease with X6 platforms is because processors have access to a steady stream of data held in-memory, so less work needs to be parallelized and handled by other processors. This means that fewer processor cores are needed to process work). Because software is priced per core, the fewer cores needed, the lower the software licensing costs.***

Other Use Cases
Beyond these four focus markets, *Clabby Analytics* also sees additional opportunities for the new X6 servers in markets where high IOPS support is needed (e.g., in data mining and database application environments that require moderate to high IOPS support; in financial services markets where there is a need to store real-time data on Flash technology such that it can be accessed quickly (speeding the decision making process and creating a competitive advantage); in surveillance and security environments where security checks are made against reference materials; in video rendering and complex CAD environments; and in multimedia streaming and video-on-demand environments where high demand requirements can be met quickly by Flash cache data delivery.

We also expect to see X6 servers deployed in the following types of environments:

- In-memory, multi-terabyte database environments that require strong resiliency;
- Virtualized environments that require access to extensive memory in order to handle I/O intensive workloads;
- Environments in which managing end-to-end performance at the server level is the goal;
- Environments that need to manage large amounts of data;
- Environments that require continuous application uptime; and,
- Environments in which removal of infrastructure complexity and cost is crucial.

The Benefits of Using X6 Servers: Faster Results, Agile Configurations, More Resiliency
IBM describes its X6 offerings and the new X6 platforms' modular design as being "*fast*, *agile*, and *resilient*." Each of these characteristics bears closer examination.

Fast
Probably the biggest benefit delivered by X6 platforms comes in processing speed. Thanks to access to large amounts of data in cache, we estimate that X6 servers will enable IBM high performance systems to process workloads approximately twice as fast as the previous generation.

Some of the reasons for this very significant increase in processing speed include:

- IBM eXFlash memory-channel storage placed on the faster memory channel (as opposed to the PCIe bus) where this memory-channel storage can serve the CPUs at higher speed (the IBM eXFlash memory-channel integrates ultra-low latency flash storage directly on the processor/memory bus for significantly better IOPs and with three times lower latency performance);
- More efficient and intelligent use of eXFlash memory-channel storage thanks to IBM FlashCache Storage Accelerator;
- The use of accelerated write drivers to deliver lower latency (the communications speed between the processor and Flash DIMMs) in a feature called WriteNow;
- The use of IBM Fast Setup and Configuration Patterns for fast provisioning; advanced memory recovery and advanced core recovery features; and many more speed and resiliency enhancements.

Agile

IBM's agility claim is backed by the ability to deploy modular books in the x3850 X6 and x3950 X6 servers in various combinations in order to add, modify or upgrade those servers easily. Books can be placed into a chassis in multiple configurations as desired (users can select the number of compute books and I/O books and the type of I/O books desired). The system can then be tailored to optimally handle specific workloads.

As the industry's first Intel Xeon-based 8-socket high-end system in a blade form factor, the Flex System X6 portfolio can also be considered more agile due to expandability options. IT buyers can start low (with 2-way configurations) and expand up to 8-way environments by using pluggable IBM Flex System Scalability Connectors (see Figure 2 –next page). When starting with a 2-way configuration IT buyers will need to purchase the appropriate Intel processors that allow for expansion to 4-way or 8-way configurations. Then, when it is time to upgrade, IT buyers make use of the scalability connectors (mentioned above) to expand systems capacity.

IBM also claims that its new X6 servers are more flexible than competitor systems due to the ability to host multiple processor technologies within the same chassis. IBM sees its X6 servers as more flexible than competing designs due to the ability to easily scale the servers to meet growth needs using 4 to 8-way rack servers as well as 2 to 8-way Flex System compute nodes.

IBM X6 platforms also offer storage flexibility. By moving data to internal eXFlash memory-channel storage, X6 servers can help significantly reduce external storage costs. Why? In this scenario, enterprises need not purchase as many storage arrays to offload data – instead, data is housed within the server. Accordingly, enterprises will need to purchase fewer physical arrays and fewer network devices – thus significantly lowering storage costs.

Finally, X6 servers introduce flexibility to the data center in the form of system footprint. The new X6 modular footprint enables IT managers to access significantly more compute power within a footprint similar to today's x86-based servers. *What this means is that an existing data center using today's x86-based servers could deliver the same amount of compute power using significantly less space*. In addition, because data moves into the server onto FlashDIMMs, fewer disk arrays (and associated networking equipment) need to be purchased and deployed. This, too, saves data center floor space. As a result of placing more computing power into scalable powerful servers – and as a result of moving data into those server environments – data center sprawl can be contained. Enterprise managers seeking to avoid the cost of building new data centers should pay particularly close attention to this observation given that building a new data center can cost almost $7 million USD according to the latest Reed Construction Data.

Figure 2 – Flex System X6 Compute Node Scalability Connectors

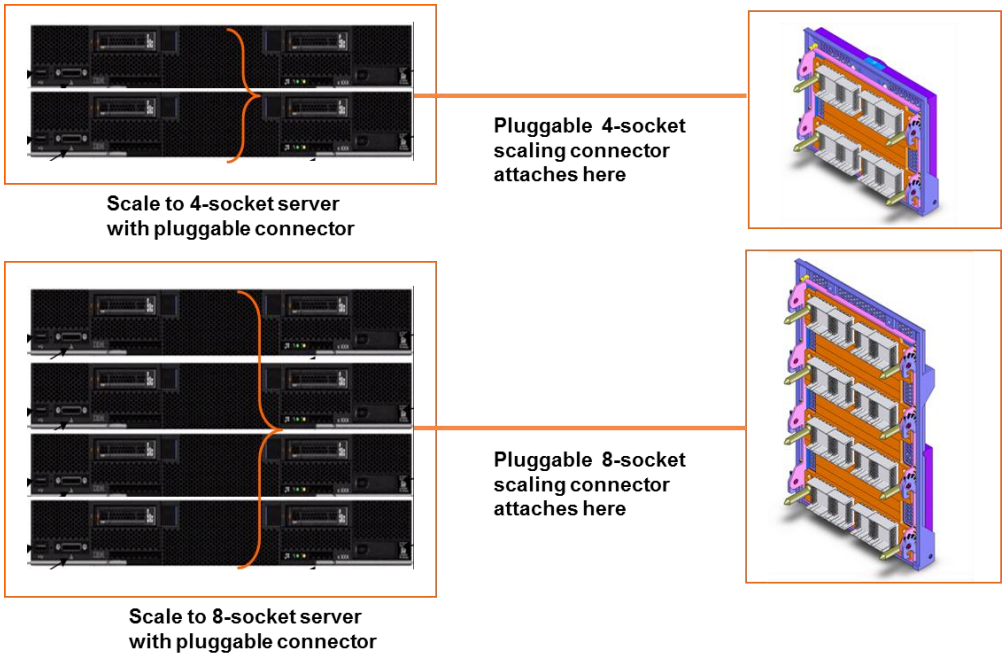

Source: IBM Corporation – May, 2014

Resiliency

One of the major design goals of IBM X6 platforms is to improve resiliency in order to maximize application uptime. To do this IBM has focused on making it easier to maintain its X6 servers (requiring fewer "touches" by IT managers and administrators), and by improving integration in virtual environments (so virtual machines are always available).

Some of the resiliency improvements found in the X6 portfolio include:

- Advanced memory recovery (this is extremely important given the amount of memory and cache that is being offered in the new X6 systems – this feature helps prevent memory failure in order to maximize uptime).
    - Monitors and identifies pages in-memory and quarantines pages when necessary to maintain uptime.
    - X6 increases system availability even with 3 times more memory than previous generations
- Self-healing architecture that proactively identifies potential failures and transparently takes necessary corrective actions. Some of these features include:
    - Advanced Page Retire that proactively protects applications from corrupt pages in memory, crucial for scaling memory to terabytes;
    - Processor High Availability that allows the platform to maintain access to networking and storage and server management during a processor failure;
    - IBM firmware advancements that maximizes platform availability by eliminating fatal error scenarios;
    - X6 modular design that reduces service time by delivering quick easy replacement of failed components.

  - Backup IMM management firmware that automatically engages if the primary fails; and,
  - Assigning work to memory components based on memory integrity (memory pages are automatically sorted based on highest memory integrity performance).
- RAID 1 mirroring for eXFlash DIMMs that drives highest availability of mission critical data placed in the 12.8TB memory channel storage.
- Virtual machine management (integrates PFA, microcode management and diagnostics into standard hypervisors through value-add modules – simplifying hardware/virtual machine integration).
  - Further, IBM's unique Upward Integration Modules (UIMs) provide unprecedented hardware visibility to the hypervisor (the program that manages virtual operating system instances on a processor) for superior system and Virtual Machine (VM) management.
  - An Upward Integration Module for standard hypervisors that enables concurrent updating of the system firmware with no impact on application performance or availability; and,
  - An Upward Integration Module for RAS that enables the creation and management of policies to maintain high availability of virtual machines.

Summary Observations

In the past, the biggest problem that has prevented x86-based servers from scaling high enough to tackle large ERP, analytics, cloud, and in-memory database workloads has been related to CPU limitations that limited the amount of memory that could be addressed. In Intel Xeon E7-8800/4800/2800 v1 processors, the amount of memory Xeon architecture could physically address topped out at between 2TB and 3TB. If the database exceeded 3TB, then the system would usually have to retrieve data from external disk – and this opened the door for I/O bottlenecks to occur – thus negatively impacting performance and throughput. With the arrival of Intel Xeon processor E7-8800/4800/2800 v2 processors, Xeon processors now have the ability to work with over 12TB of data – and this makes it possible to place new data-intensive applications on x86-based servers. IBM is the first major vendor to allow the full complement of supported memory to be accessed across both blade and rack environments by taking advantage of improvements made to the new generation Intel Xeon E7-8800/4800/2800 v2 processors that allow more memory to be accessed (The naming of the Flex nodes is meant to make it simple to know which Flex compute nodes scale to 2 sockets [x280 X6], to 4-sockets [x480 X6] or up to 8-sockets [x880 X6]).

A closer look at new IBM X6 platforms shows that IBM servers are using IBM eXFlash memory-channel storage to cache data – and that memory-channel storage is attached to the high speed memory channel to lower access latency. The eXFlash memory-channel storage is faster than competing flash technologies thanks to positioning on the faster memory channel (as opposed to the PCIe bus); this memory-channel storage can be accelerated using WriteNow; and it can be managed for performance using IBM FlashCache Storage Accelerator. As a result, data can be accessed extremely quickly; I/O bottlenecks start to disappear; and new memory-intensive applications can be deployed on IBM X6 systems.

IT executives who invest in IBM X6 servers can expect to save money by achieving higher utilization and throughput as compared to previous generation servers. The number of memory-hungry virtual machines that can be deployed on these new servers should be significantly greater than the number that could be supported on previous generation servers. We expect X6-based database servers to produce significantly higher performance than previous-generation systems

thanks to placing data in cache attached to memory channels where it can be read more quickly than having to access memory on disk using the PCIe bus. Monetary savings can also be expected in storage systems because data will reside in X6 systems – so fewer storage arrays will need to be purchased and maintained. Finally, X6 based servers can help enterprises reduce both capital and operational expenditures by enabling previous-generation servers to be consolidated, leading to higher utilization, greater performance and the ability to exploit additional capacity.

Not to be overlooked from a savings perspective are the management tools that IBM offers with its X6 servers. IBM has introduced advanced memory recovery, self-healing architecture, advanced core recovery, socket failure protection, firmware backup, virtual machine management and unique Upward Integration Modules (UIMs) to provide unprecedented hardware visibility to the hypervisor. All of these facilities help reduce the amount of human labor needed to manage X6 servers by reducing the number of "touches" required to configure, tune, and update these servers.

IT executives who wish to reduce data center costs, improve performance and throughput, improve service delivery, limit data center sprawl, reduce IT management costs – and who wish to run new memory intensive applications on x86 architecture – should look closely at new IBM X6 server environments. The performance benefits of this new architecture combined with automated management facilities offer enterprises the potential for SUBSTANTIAL TOTAL COST OF OWNERSHIP SAVINGS as well as the potential for lower labor-related management costs.

Clabby Analytics
http://www.clabbyanalytics.com
Telephone: 001 (207) 846-6662

This document was developed with IBM funding. Although the document may utilize publicly available material from various vendors, including IBM, it does not necessarily reflect the positions of such vendors on the issues addressed in this document.